

Exploiting the Cognitive and Social Benefits of Physically Large Displays

Desney S. Tan
August 2004
CMU-CS-04-154

Computer Science Department
School of Computer Science
Carnegie Mellon University
Pittsburgh, Pennsylvania 15213

Thesis Committee:
Randy Pausch (Chair)
Jessica Hodgins
Scott Hudson
Mary Czerwinski, Microsoft Research

*Submitted in partial fulfillment of the requirements for
the degree of Doctor of Philosophy*

Copyright © 2004 Desney S. Tan. All rights reserved.

This research was sponsored by various grants from Lockheed Martin, the Defense Advanced Research Projects Agency, the Office of Naval Research, and the National Science Foundation, and by National Science and Technology Board (Singapore), Agency for Science, Technology and Research (Singapore), and Microsoft Research Fellowships. The views and conclusions contained within this document are those of the author and should not be interpreted as representing the official policies or endorsements, either expressed or implied of any sponsoring party or of the U.S. or the Singapore Government.

Report Documentation Page				Form Approved OMB No. 0704-0188	
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE AUG 2004		2. REPORT TYPE		3. DATES COVERED 00-00-2004 to 00-00-2004	
4. TITLE AND SUBTITLE Exploiting the Cognitive and Social Benefits of Physically Large Displays				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Carnegie Mellon University,School of Computer Science,Pittsburgh,PA,15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES 201	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

Keywords: Large displays, physical size, cognitive and social benefits, situated cognition, Display Garden, spatial tasks, WinCuts, privacy, Spy-Resistant Keyboard, Pre-emptive Shadows.

Abstract

There exists an emerging trend in the workplace towards multiple display systems. Within these workplaces, large wall-sized displays are becoming prevalent. Although researchers have articulated qualitative benefits of large displays, little has been done to systematically quantify and exploit these benefits. My work is composed of three distinct components, each contributing to an improved understanding of physically large displays.

First, I isolate and study specific cognitive benefits unique to large displays. I present results from a series of experiments suggesting that large displays immerse users more within virtual environments and bias them into adopting egocentric strategies when performing spatial tasks. These strategies allow users to perform tasks such as 3D navigation and mental map formation more effectively on large displays than on smaller ones, even when viewed at constant visual angles.

Second, I explore social affordances offered by large displays and describe tools that I have developed to exploit these affordances. Recognizing the potential of large displays for facilitating co-located collaboration, I have developed WinCuts, an interaction technique that allows multiple users, each with their own personal computing devices, to simultaneously place and arrange information on a large shared display. I describe WinCuts as a general technique for managing information, even on standard desktop systems.

In separate work, I explore the issue of privacy on large displays. Using a novel application of an implicit memory priming paradigm, I show that people are more likely to read someone else's private content on large displays than on smaller ones, even with constant visual angles and legibility. I describe the Spy-Resistant keyboard, an interface that makes private text entry on large touch screen displays more secure against casual observers. I also present experimental results showing the effectiveness of this interface.

Finally, I explore some of the pragmatic issues surrounding the integration of large displays into our workspaces. I describe Pre-emptive Shadows, a system that uses infrared light and computer vision to eliminate blinding light cast onto an observer standing in front of a projector. I also present experimental results demonstrating detrimental effects caused by separating information within the visual field and by depth.

I close the dissertation with a summary of contributions and some future work.

Acknowledgements

I am inspired by the thoughts, the ideas, and the lives of many people. I cannot possibly create an exhaustive list, but what follows are special thanks for the people who have been the most influential both in my work and in my life.

First and foremost, I would like to thank my advisor Randy Pausch for giving me the academic freedom I needed to explore the world and find myself as a researcher, but also for stopping me on several occasions when I naïvely strayed into destructive territory. Randy has continuously provided invaluable advice that has guided not only my academic career, but also much of my life. Mary Czerwinski has also been instrumental in my development as a researcher. In her Visualization and Interaction group at Microsoft Research, Mary has provided me with a research home away from home, as well as a much welcome escape from the harsh Pittsburgh winters. She has shown me exciting ways that an effective researcher can attain impact through their work. Scott Hudson and Jessica Hodgins, the other two members of my thesis committee, have provided me with a stream of invaluable advice in shaping my research.

Dennis Proffitt at the University of Virginia has been an inspiration through his enthusiastic pursuit of scientific truth. Additionally, with his guidance, I have gained a deep appreciation for the finer points surrounding the art of experimental design. I have also immensely enjoyed interacting with members of his Perceptual Psychology Lab.

I would like to acknowledge Darren Gergle, Johnny Lee, Jeffrey Nichols, Peter Scupelli, Andrew Faulring, Adam Fass, Caitlin Kelleher, and the Stage 3 Research Group at Carnegie Mellon University, as well as Eric Horvitz, George Robertson, Brian Meyers, Patrick Baudisch, Greg Smith, and Pedram Keyani at Microsoft Research for their support and stimulating discussions on multiple and large display systems.

I thank my undergraduate advisors at the University of Notre Dame, John Uhran and Peter Kogge, for watching over me as I tried to switch majors six or seven times and helping me find my way so that I did not end up a lost, schizophrenic researcher. Robert Szczerba encouraged me in my decision to go back to graduate school. Through his example, I saw that spending over ten years in college wasn't so bad.

For having faith in my work and putting their money where their faith lay, I would like to acknowledge the National Science and Technology Board and the Agency for Science, Technology and Research in Singapore, as well as Microsoft Research for their generous fellowships that supported me through a large portion of my graduate career. Also, Lockheed Martin, the Defense Advanced Research Projects Agency, the Office of Naval Research, and the National Science Foundation have partially funded my work.

My dad Tony has given me the most important thing a man can ever give his son, an example. His strength in life has never ceased to inspire me. My mom Lorna showed me faith and taught me how to dream. It is she who instilled in me the fiery curiosity that cannot seem to be quenched. My younger sister Dawne has shown me what it means to work hard, and to keep working hard even when the rest of the world doesn't seem to care. Uncle Clement showed me the power of reason and logic in my early years, though sometimes using rather odd means like convincing me to eat tadpoles. Thanks to Ahn and Uncle Gan, who served as surrogate parents and took me under their wings when I first moved to the United States.

Finally, and perhaps most important, I would like to thank my wonderful wife Angela, for constantly encouraging, or at least tolerating, my clacking away on the computer when what she really wanted was for me to get off the machine and pay some attention to her. In good times and in bad, she has been my pillar of constancy. She is every woman I ever dreamed of and every woman I would have ever dreamed of had I known better, all in one. I love her more than anything in the world, and I'm glad beyond words that she loves me too.

Dedication

To God, through whom all things are possible

Table of Contents

Abstract	iii
Acknowledgements	iv
Dedication	vi
Table of Contents	vii
List of Figures and Tables	xi
1 Introduction	1
1.1 Motivation	1
1.1.1 Creating Environments to Support Human Cognition	1
1.1.2 Understanding the Role of Large Displays in our Environments	2
1.2 Thesis Statement.....	3
1.3 Research Approach.....	3
1.4 Research Components	5
1.4.1 Theoretical Understanding of Cognitive Benefits.....	5
1.4.2 Tools Leveraging Social Affordances.....	6
1.4.3 Pragmatics.....	7
1.5 Dissertation Organization.....	8
2 Related Work	10
2.1 Large Displays in Computing Spaces.....	10
2.1.1 Engineering Large Displays.....	10
2.1.2 Contextual Displays for Ambient Information.....	12
2.1.3 Public Surfaces for Ad Hoc Social Activity.....	14
2.1.4 Interactive Boards for Informal Group Meetings.....	15
2.1.5 Shared Displays for Collaboration.....	17
2.1.6 Large Display Environments.....	18
2.2 High-Level Effects of Large Displays.....	20
2.2.1 Arousal	21
2.2.2 Enjoyment of Content	22
2.2.3 Perception of Reality and Sense of Presence	22
2.2.4 Attention and Memory	23
2.2.5 Social Effects	24
2.2.6 Subjective Distance Preferences	24
2.3 Display Characteristics Important to Large Displays.....	25
2.3.1 Number of Pixels	26
2.3.2 Field of View	27
2.3.3 Physical Size	29
3 Large Displays Improve Performance on Spatial Tasks	32
3.1 Introduction	32
3.1.1 Exploiting User Perception to Support Cognition.....	33

3.1.2	Summary of Experiments.....	34
3.2	General Experimental Setup.....	35
3.2.1	Equipment.....	35
3.2.2	Keeping Color, Brightness, Contrast Constant	37
3.2.3	Keeping Users' Heads Still	37
3.3	Experiment 1: Physical Size Matters.....	38
3.3.1	Participants.....	38
3.3.2	Procedure	38
3.3.2.1	Guilford-Zimmerman Spatial Orientation Task.....	39
3.3.2.2	Reading Comprehension Task	40
3.3.2.3	Post-test Preference Questionnaire	40
3.3.3	Results.....	40
3.3.3.1	Spatial Orientation Task Performance	41
3.3.3.2	Reading Comprehension Performance.....	42
3.3.3.3	Preference Data	42
3.3.4	Summary	43
3.4	Experiment 2: Large Displays Bias Users into Egocentric Strategies.....	43
3.4.1	Participants.....	44
3.4.2	Procedure	44
3.4.3	Results.....	44
3.4.3.1	Effects of Strategies on Task Performance	45
3.4.3.2	Preference Data	46
3.4.4	Summary	46
3.5	Experiment 3: Large Displays Do Not Aid Exocentric Tasks.....	47
3.5.1	Participants.....	47
3.5.2	Procedure	47
3.5.2.1	Card Test.....	48
3.5.2.2	Cube Test	48
3.5.2.3	Shepard-Metzler Test.....	49
3.5.3	Results.....	50
3.5.3.1	Exocentric Task Performance	50
3.5.3.2	Preference Data	51
3.5.4	Summary	51
3.6	Experiment 4: Large Displays Improve Path Integration Performance.....	52
3.6.1	Participants.....	54
3.6.2	Task.....	55
3.6.3	Procedure	56
3.6.4	Results.....	58
3.6.4.1	Overall Task Performance	58
3.6.4.2	Systematic Component Errors.....	59
3.6.4.3	Effects of Triangle Shape.....	59
3.6.5	Summary	60
3.7	Experiment 4b: Classifying Path Integration Errors.....	61
3.7.1	Participants.....	62
3.7.2	Procedure	62
3.7.3	Results.....	63

3.7.4	Summary	63
3.8	Experiment 5: Large Displays Aid Map Formation and Memory	63
3.8.1	Participants	64
3.8.2	Procedure	65
3.8.3	Results	66
3.8.4	Summary	68
3.9	Experiment 6: Ecological Validity of Results	68
3.9.1	Participants	69
3.9.2	Procedure	69
3.9.3	Results	70
3.9.4	Summary	71
3.10	General Discussion of Experimental Results	72
3.11	Future Work	74
4	WinCuts – A Tool that Facilitates Collaboration on Large Displays	76
4.1	Initial Motivation	76
4.2	Background	78
4.3	WinCuts Interaction Technique	80
4.3.1	Basic Technique	80
4.3.2	Sharing WinCuts Across Machines	81
4.4	General Usage Scenarios	82
4.4.1	Single Machine Tasks	82
4.4.2	Multiple Machine Tasks	85
4.5	High-Level Implementation Details	86
4.6	Limitations of Current Implementation	87
4.7	Summary of Key Ideas	88
4.8	Future Work	89
5	Examining Social Issues when Working on Large Displays	91
5.1	Quantifying How Physical Size Impacts Privacy	92
5.1.1	Materials	92
5.1.2	Procedure	94
5.1.3	Results	94
5.1.4	Summary	95
5.2	Secure Password Entry on Public Touch Screen Displays	96
5.2.1	Background	97
5.2.2	Design Approach	99
5.2.3	Justification of Approach	100
5.2.4	Spy-Resistant Keyboard	100
5.2.5	Design Rationale	102
5.2.5.1	Basic Interaction Mechanism	103
5.2.5.2	Design and Layout of Tiles and Interactors	103
5.2.5.3	Specifying Shift State	104
5.2.6	User Study	104
5.2.6.1	Participants and Setup	104

5.2.6.2	Task and Procedure	105
5.2.6.3	Design	105
5.2.6.4	Results	106
5.2.7	Discussion and Future Work	108
5.2.8	Summary	110
6	Exploring Pragmatic Issues in Large Display Systems	111
6.1	Introduction	111
6.2	Pre-emptive Shadows	112
6.2.1	Implementation	113
6.2.2	Evaluation	114
6.2.3	Summary	115
6.3	Effects of Separation and Discontinuities Across Multiple Displays	115
6.3.1	Background	116
6.3.1.1	Human Vision and Peripheral Information	116
6.3.1.2	Notifications	118
6.3.2	Hypotheses	119
6.3.3	Experiment	119
6.3.3.1	Participants	119
6.3.3.2	Experiment and Setup	120
6.3.3.3	Tasks and Procedure	121
6.3.3.4	Design	122
6.3.3.5	Results and Discussion	123
6.3.4	Design Recommendations	125
6.3.5	Summary and Future Work	126
7	Conclusion	127
7.1	Summary of Work and Specific Contributions	127
7.1.1	Contributions to Theoretical Understanding of Cognitive Benefits	128
7.1.2	Contributions to Tools Leveraging Social Affordances	129
7.1.3	Contributions to Pragmatics	131
7.2	Future Work	131
	Appendix A: Selected Experiment Materials	134
	Bibliography	172

List of Figures and Tables

1 Introduction

- 1.1 In my work, I maintained constant visual angle between each of the two displays. I varied only size and the distance to the user. 3
- 1.2 The Display Garden, a rapidly configurable collection of physical display devices such as whiteboards, pin-up space, audio displays, mobile LCD panels, and large projection surfaces on various surfaces in the room..... 4

2 Related Work

- 2.1 The NCSA Display Wall-In-A-Box is a 20 projector array carefully stitched together to form a single large high-resolution display..... 11
- 2.2 Focus-in-context screens provide a large low-resolution overview of the working context around a smaller high-resolution inset of the focal information. 12
- 2.3 Kimura presents interactive montages of images on large peripheral displays to help users manage multiple working tasks. 13
- 2.4 (left) The Plasma Poster and (right) Groupcast use large displays to provide public surfaces for remote or local ad hoc social activity..... 14
- 2.5 The SmartBoard™ provides a large touch screen display that supports (left) informal group meetings or (right) presentations. 16
- 2.6 (left) Using physical tape to make conceptual sketches in car design and (right) using digital tape for similarly fluid interaction with large displays. 17
- 2.7 (left) NASA operations control room for Gemini V flight and (right) JPL war room combining individual workstations with large shared displays..... 18
- 2.8 Conceptual sketch of the Office of the Future, a spatially immersive workspace using large projection instead of standard desktop monitors. 19
- 2.9 The Stanford Interactive Room integrates a wide array of interface technologies, including a variety of large displays..... 20
- 2.10 The Dsharp widescreen display has facilitated studies exploring the effects of wide fields of view in productivity as well as spatial tasks. 29

3 Large Displays Improve Performance on Spatial Tasks

3.1 Basic experimental setup maintaining visual angles between the small and the large displays by adjusting the distance appropriately.	33
3.2 Various stages of human perception – the distal object reflects light and forms an image on the retina, which is interpreted to form the perceived image.	34
3.3 Top view schematic of the experimental setup. I maintained constant visual angles by varying display size and distance accordingly.....	36
3.4 (left) Numeric keypad input device used in the first three experiments. User working on the small (center) and large (right) displays.	38
3.5 Sample question from the Guilford-Zimmerman Spatial Orientation test. The correct answer for this question is option 5.	39
3.6 Main effect of Display Size. Users performed significantly better on the Large Display than on the Small Display.	41
3.7 Main effects of Strategy, with users performing significantly better with Egocentric Instructions than Exocentric ones. Also, results suggest that users with Unbiased Instructions perform with exocentric strategies when using the Small Display, and with egocentric strategies when using the Large Display.	45
3.8 Exocentric tests that do not benefit from egocentric strategies, or user imagining their bodies within the problem space. In each, the user has to mentally rotate images to determine if they can be of the same object.	49
3.9 Users performed no differently on any of the tasks whether using the Small or the Large Display. Egocentric strategies do not help on exocentric tasks.....	50
3.10 Results similar to Shepard Metzler (1971) experiments. The time taken for each questions is a linear relationship to the amount of rotation required.....	51
3.11 First person view of arena with the pole that users saw when performing the tasks.....	55
3.12 Diagram of terms used in the triangle completion task. Black lines represent the actual triangle; gray lines represent user response.	55
3.13 The joystick used (left); User working on the small display (center) and the large display (right).	57
3.14 Six different triangles tested in the Experiment 4.....	57

3.15	Main effects of Interactivity and Display Size. Note that users did better with Passive Viewing. There were no interactions between manipulations.....	58
3.16	(left) First person view of the world, including walls, target, and fence. (right) Map view of an example world. The user never saw this view.....	65
3.17	Similar to Experiment 4, main effects of Interactivity and Display Size. In this experiment, users benefited from Active Control. There were again no interactions between manipulations.	67
3.18	First person view of the world, which contains distinct landmarks and rich textures in Experiment 6. The target is the red flag.....	69
3.19	Main effect of Display Size. Users performed significantly better when working on the Large Display than the Small one.	71
4	WinCuts – A Tool that Facilitates Collaboration on Large Displays	
4.1	WinCuts was initially motivated by the desire to provide a tool that would allow multiple users to share information on a large display.	77
4.2	User makes two WinCuts to compare statistics between two cities. Each WinCut is an independent live region with which users can interact.	81
4.3	There are three buttons in the menubar of each WinCut: Share allows users to share WinCuts between machines, 100% returns the WinCut to original size, 1:1 returns it to original aspect ratio.	82
4.4	The most basic use of WinCuts is to specify regions of interest and get rid of all else. The user can regain scarce screen space as well as attention.	83
4.5	Users can make WinCuts of content from multiple sources, and then rearrange or even rescale them to make for easier visual comparison.....	83
4.6	Using WinCuts, users can reconfigure entire interfaces or desktops.....	84
5	Examining Social Issues when Working on Large Displays	
5.1	As with other experiments, I held the visual angle constant between the Small and Large Display conditions.	93
5.2	Users responded with target words significantly more in the Large Display condition, suggesting that they had read text more in this condition.	95
5.3	User typing with a soft keyboard on a publicly observable touch screen.	96

5.4	The Spy-Resistant Keyboard. . In the first phase of typing, the mapping phase, the user must first find the character they would like to type and note its location. For example, if they are trying to type the capital letter “Z”, they would scan the three green rows, finding the letter on the seventh tile on the second row.	101
5.5	Each Tile on the keyboard begins with a random shift state, indicated by the red underline. Hitting the Interactor moves the shift state on all Tiles. In our example, the user taps the Interactor to cycle through states and to get the underlining under the letter “Z”.	101
5.6	To complete the selection, the user drags the Interactor to the appropriate Tile. It is very difficult for an observer to reconstruct what has been typed (a letter “Z” in this case).	102
5.7	Main effect of Interface for average time to type each password.	106
5.8	Main effects of Interface for the level of security, measured by errors in guessing the password.	107
5.9	(left) We can eliminate the visual search task by laying out the keys as they normally would be. (right) The mapping is performed by blanking the keys and animating them into new positions. After mapping, the sequence of keys from left to right, top to bottom, is 9,4,8,*,2,0,5,7,1,#,3,6. To type a character, the user specifies its new position.	109
6	Exploring Pragmatic Issues in Large Display Systems	
6.1	(left) Blinding light shining on user standing between projector and wall. (right) Blinding pixels turned off using Pre-Emptive Shadows.	112
6.2	User casts IR shadow on camera lens. This matches the projected shadow and is used to turn off the appropriate pixels.	113
6.3	Luminance values for various conditions.	114
6.4	Experiment setup. I held visual angles constant between the Small and Large Display conditions. The primary display was always a small one.	120
6.5	Display conditions used in the experiment.	122
6.6	Though there were no significant differences on the Small Display for number of correct text comparisons (left), there was a significant difference between near-within and far across for the Large Display condition (right).	123
6.7	Though there were no significant differences on the Small Display for number of errors found (left), there was a significant difference between near-within and far across for the Large Display condition (right).	124

“Human brains are making the human world smarter and smarter, so that they (the brains) can be dumb in peace. Or rather, we are progressively altering our environment so that our brains ... can support intelligent choice and action by their participation in much larger networks.”

Andy Clark

Chapter 1

Introduction

1.1 Motivation

Most of us live in a mixed reality. Our daily activities are divided between the physical environment, or the real world, and a digital reality, or a virtual world. These worlds form an integral part of the way we think and act. In fact, there is reason to believe that human cognition is an intrinsically distributed phenomenon that might best be studied as a complex system occurring across individuals, physical artifacts, and symbolic representations of abstract ideas (Hollan, Hutchins, & Kirsh, 2000). Therefore, as we design mechanisms to support human-computer interaction, we must expand our focus to include these larger environments: how they augment the way we think and work, how they affect the way we interact with other people around us, and how we can best design them to create productive work environments.

1.1.1 Creating Environments to Support Human Cognition

The idea of carefully crafting the physical landscape to uniquely affect human thought and action is not a new idea. Just as a blind person's cane or a cell biologist's microscope is a central part of the way they perceive the world, so too do well-designed physical environments become integrated into the way people think and act. Kirsh (1993) classifies the functions of physical space into three categories: spatial arrangements that simplify choice, spatial arrangements that simplify perception, and spatial dynamics that simplify internal computation. He is one of many researchers who have studied not only how peo-

ple set up their workspaces to perform certain tasks, but also how they continuously manage these workspaces to facilitate thought and action.

Similarly, and perhaps more obviously, most digital environments have been built for the purpose of supporting human thought and action. For example, traditional user interface design has focused largely on the information content that lies within the virtual world. Much of this research has aimed at understanding the symbolic representation of information that most effectively communicates abstract ideas so that we can build the tools necessary for people to easily perform their tasks.

Even though we have much experience in designing both real and virtual worlds, Ishii and Ullmer (1997) observe that the two worlds remain largely disjoint and that there exists “a great divide between the worlds of bits and atoms.” In their work, they identify input devices as bridges that serve to connect the two worlds. They focus on understanding how physical objects and architectural surfaces can be used to control digital objects in the virtual world. Using their tangible interfaces, they attempt to build computing environments that support human thought and action.

However, little effort has been spent on understanding the design of the physical computer and its associated display devices (Buxton, 2001). Most work in this area has focused on pragmatic issues surrounding the changing form factors of displays, but few researchers have devoted much attention to understanding how physical affordances of these displays fundamentally affect human perception and thought. As such, design principles have been uniformly applied across a variety of display devices that offer different cognitive and social affordances.

1.1.2 Understanding the Role of Large Displays in our Environments

I assert that computer displays, which remain the dominant medium through which computers communicate information to us, also serve as bridges that connect the real and virtual worlds. Displays possess a certain duality since they exist in the real world while allowing us to peer into the virtual one. In my work, I focus my attention on user reactions to physical properties of information. Specifically, I seek to understand and exploit the affordances offered by physically large displays that exist within our workspaces.

Understanding the role of physically large displays is significant because of the emerging trend in the workplace towards multiple display systems that have the potential




Figure 1.1. In my work, I maintained a constant visual angle between each of the two displays. I varied only size and the distance to the user.

to provide abundant display space distributed throughout the environment. Such workplaces typically include several types of displays, each with characteristics that may make it more or less suitable for certain tasks. Within these workplaces, large wall-sized displays are becoming prevalent. Although researchers have previously realized that “when a display exceeds a certain size, it becomes qualitatively different” (Swaminathan & Sato, 1997), little work has been done to systematically quantify or exploit these benefits.

1.2 Thesis Statement

In my work, I seek to show that:

Information elicits fundamentally different cognitive and social reactions when presented on large wall-sized displays as compared to smaller displays, even at identical visual angles (see Figure 1.1). These reactions can be quantified and understood in controlled experiments and can be exploited to design display systems that make users more productive than they were on traditional systems.

1.3 Research Approach

I have tried to be opportunistic in directing and shaping my research. I believe that there is a fine balance to be struck between staying on the path towards a larger vision, and exploring sometimes tangential but often interesting problems that inevitably arise along that path. The former ensures that we do not get lost as researchers, wandering the design space for inconsequential problems to solve. The latter ensures that we do not become so



Figure 1.2. The Display Garden, a rapidly configurable collection of physical display devices such as whiteboards, pin-up space, audio displays, mobile LCD panels, and large projection surfaces on various surfaces in the room.

engrossed in a single problem that we lose sight of the greater goals and vision that drive our higher level agenda.

Much of my dissertation work has been motivated by a combination of anecdotal evidence, informal observations, and established theoretical work in psychology, human-computer interaction, and computer science. It is through these channels that I was able to identify and focus on the areas in which I thought large displays were likely to have the most impact on user performance.

My general approach was a three-pronged strategy including: (1) combining theoretical work with empirical evidence to identify display characteristics most likely to impact the way we think and work; (2) designing controlled experiments to isolate and understand effects more completely; and (3) deriving design principles and building real-world systems that make users more productive.

To explore the issues surrounding the integration of physically large displays into the workspace, I created a system called the Display Garden. This system is a rapidly configurable collection of physical display devices such as whiteboards and pin-up space, audio displays, mobile LCD panels, and large projection displays on various surfaces in the room (see Figure 1.2).

Although I do not view the creation of this system as a significant intellectual contribution in and of itself, working within the Display Garden during the course of my work has provided me with a deep appreciation for the nature of physically large displays. I

believe that working within any new system is integral to a complete understanding and appreciation for the subtleties presented by the system. Ideas garnered from working within the Display Garden form the basis of much of my work in this dissertation. Additionally, the Display Garden has provided the hardware infrastructure that has allowed me to rapidly prototype new ideas. In this role, it has served as a tool for me to apply and validate design principles I have formulated for building information environments with large displays.

1.4 Research Components

My work is composed of three components: (1) a theoretical understanding of cognitive benefits of large displays; (2) tools and interface techniques leveraging social affordances offered by large displays; and (3) examination of some of the pragmatic issues surrounding the creation of these display environments. I describe these in more detail in the following sections.

1.4.1 Theoretical Understanding of Cognitive Benefits

First, we must isolate and study characteristics unique to large displays so that we form a theoretical basis for understanding how they affect the way we think and work. I believe that taking a bottom up approach and understanding each of these fundamental characteristics in isolation rather than taking a top down approach and studying how one display technology differs from another will be much more productive in the long run. Such an approach will allow us to build a general theory that explains effects induced by various display technologies simply by recombining our understanding of display characteristics and then studying their specific interactions.

Researchers have already begun to isolate certain interesting characteristics of large displays. However, most researchers have assumed that larger displays fill a greater percentage of the viewer's visual field, and physical size is often confused with visual angle, or field of view. In fact, while researchers have studied the effects of display characteristics such as field of view, resolution, brightness, contrast, and color, little has been done to systematically isolate the effects that physical size has on the way users react to information. In my work, I aim to develop a theoretical understanding of physical display size as it relates to the way we think and work.

One of the areas I have explored in detail is spatial cognition. In this work, I show that physical display size, independent of other display characteristics, affects the way we think and work. In fact, physically large displays seem to improve performance on many spatial tasks, even when I held factors such as field of view constant. Using a series of such tasks, I show that this effect can most likely be attributed to physical size of the displays inducing users into adopting different cognitive strategies. In fact, small displays seem to bias users into using exocentric strategies and large displays seem to immerse users more within the virtual environment and bias them into more efficient egocentric strategies.

1.4.2 Tools Leveraging Social Affordances

Second, I have used physically large displays as a means to motivate thought about issues that did not exist when using traditional desktop displays.

Specifically, thinking about using large displays in the environment for collaboration has made me consider the scarcity of screen space in our everyday computing systems. While the screen space problem did previously exist on traditional desktop systems, large displays exaggerated the problem and solving it led to an interesting solution for desktop computing in general. As a solution to this problem, I have designed an interaction technique that allows users to replicate arbitrary regions of existing windows into independent windows called WinCuts. WinCuts may either be used on the same machine as the source window or be shared with remote machines. Each WinCut is a live view of a region of the source window with which users can interact. By allowing users to choose exactly what is shown and where, this technique allows users to easily and effectively manage their information and screen space.

Another observation when working on physically large displays is that a certain amount of information privacy is lost. A common explanation for this loss in privacy is the higher legibility of information presented on large displays. Because large displays are typically viewed from a distance that is not proportionally scaled with the increase in display size, they often provide a larger visual angle, making them easier to see and read. While I agree that this contributes to the loss of privacy, I assert that there are more subtle social cues that may also contribute to this effect. Using a novel application of an implicit memory priming paradigm to measure whether or not someone has read a particular pas-

sage of text, I show that people are more likely to read text presented on a larger display, even when visual angles and legibility are held constant.

Since users have intrinsically less privacy when working on large displays, it is difficult to perform certain actions, such as entering private text passwords, without being observed. This is especially true with large touch screen displays on which both the interaction and the result of the interaction are visible. Someone watching the typist interact with an onscreen soft keyboard on a touch screen display can fairly easily reconstruct text that has been entered, an activity commonly known as shoulder surfing. I have devised a novel approach to designing keyboards for entering private text on public touch screen displays. I describe one instantiation of such an interface, called the Spy-Resistant Keyboard, and present evaluation of its effectiveness.

1.4.3 Pragmatics

Finally, I have examined some of the issues surrounding deploying large displays in our workspaces and creating environments consisting of a myriad of display devices. The task of integrating all our understanding and tools to create rich computing environments with multiple display systems is intrinsically an engineering effort. As with any other sizable engineering effort, we can expect to encounter technical problems, some tied to particular technologies, but others more universal in nature.

For example, since I was front projecting to create large displays in the Display Garden, users in the room often found themselves working between the projector and the display surface. This caused undesirable projection on the user as well as temporary blindness from looking into the bright light of the projector. To alleviate this problem, I have developed Pre-emptive Shadows, a technique that uses an infra-red camera-projector system to detect and turn off pixels that would otherwise be needlessly cast upon users' bodies and faces.

Furthermore, having multiple displays in the Display Garden enlarges the physical display area, allowing the system to present information across much wider visual angles from the user. Since displays are placed at different depths or framed by physical bezels, physical discontinuities are also introduced into the presentation of information in the workspace. Relatively little is known about the how to best present information to the user given these display characteristics. In my work, I use a divided attention paradigm to

explore the effects of visual separation and physical discontinuities when distributing information across multiple displays. Results show reliable, though relatively small, detrimental effects when information is separated within the visual field, but only when coupled with an offset in depth. Surprisingly, physical discontinuities such as monitor bezels and even separation in depth alone do not seem to affect performance on the set of tasks tested. This has implications for industrial design of multiple display systems, which are quickly becoming commonplace.

1.5 Dissertation Organization

In this chapter, I have briefly presented my high level motivation and goals as well as the approach I took in performing much of my work with large displays.

In Chapter 2, I discuss related work that has contributed to a better overall understanding of large displays and how they can be used in our computing environments. The work presented within this chapter forms the foundation for much of my work in this dissertation.

The core of the dissertation is broken into four chapters. Although all my work was motivated by thinking about large displays and how we can best use them to design environments that support human thought and action, each of these chapters is a fairly distinct set of work.

In Chapter 3, I present a series of experiments showing that physically large displays, even when viewed at identical visual angles to smaller ones, affect the way we perceive certain information and can increase task performance on spatial tasks.

In Chapter 4, I describe the WinCuts interaction technique, a tool designed and built to support co-located collaboration on large displays. I discuss how WinCuts serves as a much more general technique for managing information across various tasks and computing environments.

In Chapter 5, I explore social affordances of large displays, especially with regard to information privacy. I also present the Spy-Resistant Keyboard, an interaction technique designed to allow users to enter private text such as passwords even when they are being carefully observed on large public touch screen displays.

In Chapter 6, I explore some of the pragmatic issues surrounding the use of large displays in our environments. Specifically, I present Pre-emptive Shadows, a technique that ensures light from front projection screens do not blind users standing in front of them. Also, I explore the effects that distributing information on multiple displays throughout the environment has on task performance.

Finally, in Chapter 7, I summarize the work and contributions presented in this dissertation. I also discuss directions for future work.

I include key materials from the main experiments in this dissertation within an appendix that follows.

While much of the present text is new material, a few sections draw on content from previously published articles, namely sections 3.3 (Tan, Gergle, Scupelli, & Pausch, 2003), 3.6 (Tan, Gergle, Scupelli, & Pausch, 2004), 5.1 (Tan & Czerwinski, 2003a), 6.2 (Tan & Pausch, 2002), and 6.3 (Tan & Czerwinski, 2003b), as well as parts of Chapter 4 (Tan, Meyers, & Czerwinski, 2004).

*“The moments of the past do not remain still;
they retain in our memory the motion which drew them towards the future,
towards a future which has itself become the past,
and draw us on in their train.”*

Marcel Proust

Chapter 2

Related Work

This chapter, broken into three sections, provides a general overview of the state of research conducted around large displays. In the first section, I describe projects that have utilized large displays in one way or another. In the second section, I examine work that has explicitly studied the high level effects of using these large displays as a whole. And in the third section, I examine work that has studied the effects of individual characteristics unique to large displays. More specific related work is distributed throughout the dissertation as appropriate.

2.1 Large Displays in Computing Spaces

Large displays have been used extensively in a variety of projects and scenarios. While it is beyond the current scope to exhaustively document every project that has ever used a large display, this section highlights some of the work that has explicitly revolved around large displays, or that makes interesting use of such displays.

2.1.1 Engineering Large Displays

We are at a point in time when technology trends and user demands are fueling the display industry to produce larger and larger desktop displays for less and less money. However, for a variety of reasons, high resolution wall-sized displays remain fairly expensive. Hence, the engineering challenge of building these displays out of commodity parts has attracted the attention of several groups. Many of these groups have focused on scalable



Figure 2.1. The NCSA Display Wall-In-A-Box is a 20 projector array carefully stitched together to form a single large high-resolution display.

rendering for large displays (e.g. Humphreys & Hanrahan, 1999; Li et al., 2000), creating complex graphics architectures necessary for composing high-resolution images useful in many domains, such as data visualization. For example, Schikore et al. (2000) have developed a system that displays up to 15 times the number of pixels on a typical desktop display so that the Department of Energy can visualize complex data sets.

Others researchers have focused their efforts on the hardware associated with the actual displays. In desktop computing, researchers have explored the use of multiple display systems (Dunn & Mikes, 2001), claiming a growing trend for users to have multiple monitors associated with their desktop machines. Beyond the desktop, many researchers have worked on combining multiple desktop or projection displays to form large tiled display walls. The PowerWall and InfinityWall (Czernuszenko, Pape, Sandin, DeFanti, Dawe, & Brown, 1997), as well as the National Computation Science Alliance Display Wall-In-A-Box (see Figure 2.1) are examples of such systems. A smaller portion of this work has involved less standard display form factors, such as curved or domed displays (e.g. Raskar, van Baar, Willwacher, & Rao, 2004).

Extending this work, researchers have combined multiple display walls to form spatially immersive displays. Spatially immersive displays are systems that surround the viewer with a panorama of imagery (Bryson, Zeltzer, Bolas, de La Chapelle, & Bennett, 1997). These displays are typically room-sized and accommodate several viewers. Probably the best known spatially immersive display is the Cave Automated Virtual Environment (Cruz-Neira, Sandin, & DeFanti, 1993), usually a room composed of up to 6 large displays, optimally one for each of the four walls, the floor, and the ceiling. Researchers

have explored a wide range of techniques for improving the user experience in such systems, including stereoscopic viewing and seam elimination (Schell & Shochet, 2001).

Because of the rate that display technologies are evolving, I have kept my work fairly well divorced from specific technologies. Instead, I have studied fundamental psychophysical phenomena that cause us to react to various display characteristics. In this way, I not only contribute a better understanding of human cognition to the field of psychology, but also derive more general principles for designing and building display systems.

2.1.2 Contextual Displays for Ambient Information

Because large displays are intrinsically more visible than smaller ones, they can be placed further away or off in the periphery of human vision without making content harder to see or read. Recognizing this capability, researchers have explored the use of large displays to unobtrusively provide contextual information that could be useful to users as they perform their focal tasks on more traditional displays.

In the Prairie system, designed to utilize large displays for distributed knowledge management and collaboration, Swaminathan and Sato (1997) identify and support at least four distinct types of contextual information: (1) organizational context, the relationship of a community of users to other communities; (2) social context, the social activities in a community such as presence and current task; (3) work context, how various work objects on the display are related to each other; and (4) navigational context, the path through which a user reaches a particular object.




Figure 2.2. Focus-in-context screens provide a large low-resolution overview of the working context around a smaller high-resolution inset of the focal information.



Figure 2.3. Kimura presents interactive montages of images on large peripheral displays to help users manage multiple working tasks.

Baudisch, Good, Bellotti, and Schraedley (2002) provide a large low-resolution overview of the working context around a smaller high-resolution focal screen (see Figure 2.2). In a series of experiments, they showed that the persistent presence of contextual information made users more efficient at tasks that required them to view the focal information at multiple levels of detail. MacIntyre, Mynatt, Volda, Hansen, Tullio, and Corso (2001), in their Kimura office environment, assist users to manage multiple working tasks by presenting interactive montages of images on large peripheral displays (see Figure 2.3). These montages not only remind users of past actions, but also serve as contextual cues into pending tasks and projects.

Tan, Stefanucci, Proffitt, and Pausch (2001) build on the principle that the contextual information we incidentally encode when we acquire information in the real world serve as strong memory cues for later retrieving that information. In their InfoCockpits system, they utilize large peripheral projection displays to show different scenes of distinct ‘places.’ These places provide the context that serve as cues to remember more information. They show a 56% improvement in memory for information presented with the InfoCockpit system as compared to a standard desktop display system. They hypothesize that the greater the sense of presence invoked by the display, the better the memory for learned information. They do not, however, explicitly explore how the displays or their physical size affect this sense of presence.

Large displays have also been used to provide ambient non-information bearing content. For example, in the BlueSpace workplace (Lai, Levas, Chou, Pinhanez, & Viveros, 2002), large displays not being used for focal tasks automatically project artwork or mimic windows by displaying images from outdoors webcams.



Figure 2.4. (left) The Plasma Poster and (right) Groupcast use large displays to provide public surfaces for remote or local ad hoc social activity.

2.1.3 Public Surfaces for Ad Hoc Social Activity

In addition to providing ambient information to individuals, researchers have explored placing large displays in key locations such as outside office doors in order to broadcast ambient content that supports social activities in public spaces (Fitton & Cheverst, 2003; Moran, Saund, van Melle, Gujar, Fishkin, & Harrison, 1999; Russell, Trimble, & Wales, 2002). Many of these serve as digital bulletin boards that allow their owners to display text, static images, or rich media content (e.g. Fass, Forlizzi, & Pausch).

As an extension to the bulletin board metaphor, Churchill, Nelson, Denoue, and Girgensohn (2003) connect multiple large display interactive bulletin boards called Plasma Posters across the network, allowing board owners to post content to multiple locations at once (see Figure 2.4). Similarly, Grasso, Roulland, Snowdon, and Muehlenbrock (2002) built a large display system called the Community Wall to support information sharing and discovery across communities of practice.

McCarthy, Costa, and Liongosari (2001) explore the use of peripheral displays in three workplace contexts: within an individual office, directly outside the office, and in a common area. On the office displays, they present content useful to the individual. Within the other two contexts, they explore the kinds of information that users would like to share with passersby, as well as interaction mechanisms that could aid informal conversations between either local or remote viewers of the displays. While they do not explicitly study the effects of the displays themselves, it is interesting that they chose to use smaller displays for personal viewing and much larger ones as their public displays.

Greenberg and Rounding (2001) built Notification Collage, a system that allows users to post text notes and other media and to converse via live video and audio feeds. Using this system, they explored how personal peripheral displays as well as large public displays could enhance casual interaction and communication between users in a community. In their Dynamo system, Izadi, Brignull, Rodden, Rogers, and Underwood (2003) allow sharing and exchange of information across public displays to support opportunistic meetings in public settings. The large display, in this project, serves as a shared digital and physical space on which users can collaborate.

Brignull and Rogers (2003) focus on understanding how groups of people socialize around large displays. Using their Opinionizer system to study public interaction flow around large displays, they present suggestions for designing public displays that get users' focal attention as well as encourage users to interact with the display and with others around them. Extending this work, Streitz, Röcker, Prante, Stenzel, and van Alphen (2003) use three different distance-defined zones to define the semantics of interaction around their large display called the GossipWall. In the furthest, the ambient zone, the large display shows general information. As the user moves toward the display and into the notification zone, the system shows information relevant to the individual or the group surrounding the display. Finally, a user in the closest zone, the interaction zone, can interact with information on the display either by directly touching the display or by using a variety of mobile devices.

In their book, *Public and Situated Displays*, O'Hara, Perry, Churchill, and Russell (2003) provide a more complete overview of this area of research.

2.1.4 Interactive Boards for Informal Group Meetings

Large displays have also been used to support small informal group meetings. In this capacity, large displays have been used as central displays or drawing surfaces that allow easy presentation and capture of ideas. Recent work on computer-supported meeting environments has recognized the importance of these central display surfaces. Meeting rooms such as Colab (Stefik, Foster, Bobrow, Kahn, Lanning, & Suchman, 1987), Capture Lab (Mantei, 1988), and Project Nick (Cook, Ellis, Graf, Rein, & Smith, 1987) all utilize one or more large displays as a major focus of group work. In fact, Mandryk, Scott, and Inkpen (2002) articulate this principle when they identify display size as an important factor in comparing collaborative systems.



Figure 2.5. The SMART Board™ provides a large touch screen display that supports (left) informal group meetings or (right) presentations.

The Liveboard system (Elrod et al., 1992) uses a directly interactive, stylus based, large area display to complement other personal computing devices. It also provides a large shared workspace around which groups can collaborate. Initially, Liveboard simulated a standard whiteboard by allowing freehand drawing and erasing. However, applications such as Tivoli (Pedersen, McCall, Moran, & Halasz, 1993) have extended that to include sorting, categorizing, and annotating functionality that takes advantage of the computational power offered by the new medium. Interestingly, Liveboard-like systems are now commercially available (e.g. SMART Board™, see Figure 2.5). In the Flatland project, Mynatt, Igarashi, Edwards, and LaMarca (1999) explored supporting long-term informal use of such interactive large display systems in individual office settings. They further extend whiteboard functionality by interpreting high level content semantics rather than operating on lower level strokes.

With the emergence of new functionality comes the need for better sensing technologies that support interaction around large screen devices. Much research attempts to address this need (e.g. Deitz & Leigh, 2001; Leibe et al., 2000; Matsushita & Rekimoto, 1997). There has also been work to improve interaction techniques associated with these devices. For example, researchers have investigated interaction techniques which facilitate working across multiple pen-based devices (Rekimoto, 1998; Rekimoto & Saitoh, 1999). In this work, they demonstrate their pick-and-drop technique as a useful mechanism for users to use multiple personal tablets with a shared whiteboard in a collaborative setting. Nakagawa, Oguni, Yoshino, Horiba, and Sawada (1996) discuss user interfaces for large screen displays and propose GUI widgets for such environments. Balakrishnan, Fitzmaurice, Kurtenbach, and Buxton (1999) analyze physical tape drawing in the auto-




Figure 2.6. (left) Using physical tape to make conceptual sketches in car design and (right) using digital tape for similarly fluid interaction with large displays.

motive industry and describe digital tape, an interaction technique designed for large displays (see Figure 2.6). Guimbretière (2002), in his dissertation, explores tools and techniques for the fluid interaction with large display surfaces. Among others, he discusses a new menu system called the FlowMenu that allows the user to execute a wide variety of actions in a single pen stroke, and PostBrainstorm, a tool for easily organizing sketches and other information.

2.1.5 Shared Displays for Collaboration

In addition to facilitating ad hoc social interaction and small informal meetings, large displays have also been used to support more formal cooperative work needed for operating on much larger amounts of information than a single person can handle. In fact, many researchers have built systems that use large public displays to support focused, time-critical collaboration. Such systems can already be seen in control rooms of complex real-world systems such as industrial plants or in large planning scenarios such as in military command rooms (see Figure 2.7).

For example, the MERBoard system (Trimble, Wales, & Gossweiler, 2002) was designed to help NASA scientists analyze data from Mars rovers, and the eWhiteBoard system (Bercowicz, Barnett, & Chueh, 1999) supports scheduling in a cardiac catheterization center. In their work, Dudfield, Macklin, Fearnley, Simpson, and Hall (2001) explore the use of panoramic displays to facilitate shared mental models of information, as well as to improve situation awareness and team decision making. In their studies, they found strong user preference for shared large displays, with users reporting improved situation awareness and decision making. However, quantitative analysis of objective data did not support this preference. They hypothesized that this disparity might have been due to the



Figure 2.7. (left) NASA operations control room for Gemini V flight and (right) JPL war room combining individual workstations with large shared displays.

lack of experimental control over the simulation or to insensitivity of objective measures. Other studies that have also concentrated on realistic scenarios in similar military settings support their preference data suggesting that teams do indeed perform better when working on shared large displays (Hiniker, 1998; Hiniker & Entin, 1992).

The Courtyard system (Tani, Horita, Yamaashi, & Tanikoshi, 1994) was built to support coordination and division of labor by integrating an overview on a shared large display with detailed views on individual personal displays. Courtyard allows users to access per-user detailed information on their individual screens simply by moving their mouse pointer off their individual screen and pointing to an object on the shared screen.

Recently, Mark (2002) has described the idea of Extreme Collaboration, work performed within technology ‘war rooms’ consisting of individual workstations clustered around public displays and other shared resources. She presents a case study of such a war room used at the NASA Jet Propulsion Laboratory and examines issues surrounding working in such an environment (see Figure 2.7). She concludes that a delicate balance must be struck between electronic and social networks in order to optimize the flow of information.

2.1.6 Large Display Environments

In addition to integrating large displays into more traditional computing or meeting environments, many researchers are creating entirely new computing environments built around large displays. For example, in the Office of the Future, Raskar, Welch, Cutts, Lake, Stesin, and Fuchs (1998) create spatially immersive displays by projecting high-resolution graphics and text onto all the objects and surfaces in a room (see Figure 2.8). In following work, Welch, Fuchs, Raskar, Towles, and Brown (2000) highlight the im-



Figure 2.8. Conceptual sketch of the Office of the Future, a spatially immersive workspace using large projection instead of standard desktop monitors.

portance of the physical scale and high-resolution imagery, which they claim supports a more natural means of interacting with data.

In contrast to these sophisticated, technologically advanced offices, Bishop and Welch (2000) have created a simplified prototype Office of the “Real Soon Now”. One of the key innovations in this office is that they use large screen projection instead of conventional monitors. Using such displays, they claim improved social and technical interaction, better ergonomics, as well as higher information content. They also discuss issues they had with heat, noise, and brightness from the projectors, getting rid of seams with multiple projection displays, as well as privacy and cost concerns.

The Stanford Interactive Room project (see Figure 2.9) provides a wide array of interface technologies, utilization of distributed computation power, as well as highly flexible infrastructure that allows incremental addition and use of new technologies (Johnson, Fox, & Winograd, 2002). Researchers in this project experiment with multi-device, multi-user environments to facilitate fluid group interactions. One component of the interactive room is the Interactive Mural, a large, high-resolution, tiled display constructed using eight overlapping projectors driven by multiple independent graphics accelerators (Guimbretière, 2002).

In the i-LAND project, Streitz et al. (1999) create an environment that supports cooperative work of dynamic teams with rapidly changing needs. i-LAND consists of several ‘roomware’ components, or computer-augmented objects that integrate physical elements with digital information technology. The DynaWall is a large touch sensitive display on



Figure 2.9. The Stanford Interactive Room integrates a wide array of interface technologies, including a variety of large displays.

which multiple users can directly interact. Users can also place information from each of the other roomware elements onto the DynaWall for easy shared viewing. New techniques, such as “take and put” and “shuffle” were developed for users to comfortably interact with the physically large DynaWall.

Finally, researchers have used large displays to create immersive rooms for other reasons. For example, Bobick et al. (1999) use large displays, coupled with physical objects and digital sensing mechanisms, to create an interactive narrative playspaces for children in their KidsRoom.

2.2 High-Level Effects of Large Displays

Most of the work described in the previous section had to do with engineering large displays or large display systems within various scenarios and environments. In this section, I describe high level behavioral responses to large displays. This work, done mainly in the field of media communications, explores the high-level impact of screen size on viewer responses to media content.

Screen size has been a critical feature of film presentations since the transition from Kinescope and Mutoscope peep shows to projected images (Belton, 1992). The film industry recognized early on that the power of the film image could be attributed at least in part to physical size and promoted the new technology as being superior to live stage performances because it made things “larger than life” (Verdac, 1968).

While possible effects have been discussed for a long time, effects induced by physical size are only beginning to be investigated empirically. The research on screen size in media communications is motivated by two concerns: (1) crafting media that scales across an increasingly diverse set of audiovisual displays; and (2) a theoretical interest behind understanding whether and to what extent screen size causes media users to experience a sense of presence (Grabe, Lombard, Reich, Bracken, & Ditton, 1999). Although little of this work directly explores effects of large displays on productivity tasks, the understanding gained in these studies can be broadly applied across a series of tasks and serve as a good starting point for my current research.

2.2.1 Arousal

In his bio-informational theory, Lang (1995) regards emotions as motivationally tuned states of readiness that are products of Darwinian evolution. The theory suggests that two primary motive systems underlie all affective responses, the appetitive system, which causes approach emotions, and the aversive system, which causes avoid emotions. In this model, fundamental attributes of stimuli such as size, color, and motion, are seen to be influential to emotional response. Building upon this theory, Detenber and Reeves (1996) examine the effects of motion and image size on the emotional response of viewers. They showed subjects images from television and film selected to elicit a series of emotions. Using a self-report measure of emotional response called the Self-Assessment Manikin (Lang, 1980) they found that subjects felt much stronger emotions when viewing content on a large 70" display as opposed to a smaller 35" one. They concluded that the form of the message is just as important to understand as the actual symbolic content. Similarly, Lombard, Reich, Grabe, Bracken, and Ditton (2000) performed a similar comparison, but used skin conductance and other physiological measures to measure arousal. They too found significant differences, with greater arousal occurring in large display conditions.

In their work, Reeves and Thorson (1986) asserted that image size affects sensory processing more than semantic processing. They claimed that image size affects arousal ratings, but do not change evaluations of valence. That is, a sad scene will evoke a sad response whether seen on a small or a large display. The difference will be in the magnitude of sadness experienced by the viewer. Interesting effects on productivity tasks, then, should occur in situations where it is important for magnitude of arousal to be generally higher. Ideally, higher arousal will cascade into higher level cognitive effects that can increase task performance.

2.2.2 Enjoyment of Content

Many researchers have explored the relationship between screen size and viewer enjoyment using different forms of media content. Lombard, Reich, Grabe, Bracken, and Ditton (2000) showed viewers reporting a greater sense of enjoyment when watching rapid point of view movements on large screens. Similarly, Ohtani and Mitshuhashi (1973) showed that users preferred larger (20", 42", and 70") television displays when watching dramatic content. However, they also showed that large displays increased dizziness and fatigue when users viewed scenes of a fast paced horse race.

Contrary to these results, Kim (1996) reported that image size had no effect on reported "liking" of an infomercial for home exercise equipment. Detenber and Reeves (1996) found screen size had no influence on the "pleasantness" of still or full motion images taken from television and film. Lombard, Ditton, Grabe, and Reich (1997) found no effect for screen size on reported viewer enjoyment of television content taken from a variety of genres. In fact, they reported that viewers preferred the small screen.

It seems that viewer enjoyment and preference are dependent on a fairly complex relationship between screen size and specific genres of media content. In my work, I am not particularly concerned with understanding how each type of media content interacts with screen size. However, these results suggest that we must be careful to validate effects of physical display size across a series of tasks as they may be sensitive to small manipulations in the nature of content shown.

2.2.3 Perception of Reality and Sense of Presence

There is substantial evidence supporting the idea that larger displays promote higher levels of perceived reality and a greater sense of presence within content viewed. Larger displays with wider fields of view fill more of the visual field, or occupy more of the peripheral vision. This means that the boundary between the screen and rest of the environment is farther in the corners of vision. This makes the boundaries less noticeable, which has important implications for responses such as arousal and presence (Reeves & Nass, 1996). A technical report on big-screen televisions found that increasing visual angle through large image sizes and closer viewing distances led to greater subjective evaluation of the sensation of reality (Hatada, Sakata, & Kusaka, 1980).

Neuman (1990) found that for high-resolution images, a large 180" display yielded reports of "dramatically increased sense of realism" over a smaller 35" display. Reeves, Detenber, and Steuer (1993) used scenes from four action-adventure entertainment films and found that subjects reported feeling more "a part of the action" when using the large 70" screen as opposed to the 35" one. Subjects also reported that clips were more realistic on the large display.

Similar reports exist in the virtual environment and simulation literature. Bystrom, Barfield, and Hendrix (1999) assert that the more inclusive, extensive, surrounding, and vivid the display, the higher the potential of presence. In fact, when users are present in virtual environments, the location of their physical bodies are often construed as being contained within that space rather than looking at it from the outside. They hypothesized that it is in this state that users are most effective in these environments.

2.2.4 Attention and Memory

Another effect that has been extensively explored is how physical display size affects attention and memory for content. Most findings suggest that large displays induce heightened levels of both attention and memory. For example, using stimuli derived from popular movies, Reeves and Nass (1996) found that larger 90" pictures were more arousing, better liked, and better remembered than their smaller counterparts, shown on 22" screens. Detenber and Reeves (1996) found that subjects who watched images on a large 90" screen remembered more images directly after viewing them than subjects who viewed them on a 22" screen. Reeves, Lombard, and Melwani (1992) found that people appearing on a large 68" screen were given greater attention, people seen from a closer viewing distance were remembered better, and people shown in close up shots were given less attention but were remembered better. Also, de-Brujin, de-Mul, and van-Oostendorp (1992) examined the impact of screen size on learning text. They showed that subjects who learned from a 15" computer monitor learned material more quickly than those reading from a 12" screen.

Kim (1996) found that larger images led to greater sensitivity for facts presented in a 15 minute infomercial. Interestingly, Kim's results additionally suggest that the impact of changes in image size was greater when ambient light was present so that other parts of the viewing environment were clearly visible. This indicates that environmental context is important to the way we perceive physical size. Unfortunately, there are many cues that

allow us to perceive physical size, and little has been done to isolate the contributions and interactions between each of these.

However, in separate work, Reeves, Detenber, and Steuer (1993) found results directly opposed to these, that large 70" screens generated less attention and memory for people. They concluded that the larger screens provided compelling visual and auditory experiences that were exciting and well-liked, but that caused distractions and were not conducive to mindful processing of information. I believe this indicates that we must be careful in designing large display systems and content. Even though large displays generally evoke a greater level of attention and memory, this attention and memory could be easily misdirected, adding to cognitive load and leaving fewer mental cycles for processing intended content.

2.2.5 Social Effects

While there has not been a large amount of effort devoted directly to studying the effects that display size has on social effects such as trust and respect, researchers have attempted to understand the effects that proximity has on the way people interact with each other. In his landmark work, Hall (1966) developed the field of proxemics, the study of human use of space within the context of culture. Following in this tradition, researchers have used physical size to indirectly control perceptions of proximity when content is mediated by various media. For example, in his dissertation, Grayson (2000) used a wide range of approaches to demonstrate that perceptions of proximity do exist in video-mediated environments and that they do lead to differences in communication behavior. He found that when remote collaborators appeared nearer, either because the display was physically closer, or because the remote participant was larger on the display, users tended to be more interactive. While other researchers have informally speculated on other possible effects, little work has been done to formally articulate or quantify them.

2.2.6 Subjective Distance Preferences

Media communication researchers have evaluated viewing distance preferences as a function of various display factors, including image size and resolution. Results for preference of viewing distance for various sizes of screens are varied. Some researchers (Duncanson & Williams, 1973; Lombard, 1995; Nathan, Anderson, Field, & Collins, 1985) show that larger screens make viewers choose a proportionally greater distance from which to view the images. Others (e.g. Jesty, 1958; Westerink & Roufs, 1989) have found the ratio to be

constant regardless of image size. Yet others (Ribchester, 1958) have found that the ratio decreases with increasing display size.

One explanation for these results could be that the ratio of preferred viewing distance to image height is a constant that varies based on other factors such as image resolution. For example, Fish and Judd (1991) showed that for standard NTSC video, the preferred viewing distance to height ratio is about 7. However, Lund (1993) tested with images ranging from 11 to 123 inches and showed in a series of five experiments that the ratio of viewing distance to image height decreased from 7.4 to 3.1. He argues that contrary to predictions, the ratio actually decreases as image size increases, that ratios are relatively unaffected by resolution, and that the non-linear curve is due to viewers selecting their viewing positions to optimize a sense of presence or reality. Regardless of the findings, all these studies propose different hypotheses to explain the different viewing distance preferences. However, none have concretely proven or disproven these hypotheses.

These results suggest that the preferred viewing distances are probably only partially a function of the size of the display. In fact, other factors such as quality of picture and specific content may also play a part in preferences. Because many of these experiments do not fully describe their experimental setups and the particular technologies used in tests, it is hard to make definitive comparisons. However, most of the data seems to fit relatively well with recommendations suggesting a distance to height ratio of 10 for personal televisions (Sadashige & Ando, 1984) and 2 to 4 for large theaters (Kaufman & Christensen, 1987).

2.3 Display Characteristics Important to Large Displays

As described in the previous sections, there has been a large amount of work done both in constructing large display systems as well as understanding high level effects of large displays, especially as they apply to media communications. Unfortunately, results presented in most of these studies have several limitations. First, although researchers have recognized that large displays are an integral part of their display systems, they have seldom explicitly studied the specific affordances and effects of these displays. As such, we can only draw qualitative conclusions, and only for the set of displays that have been used. Second, in many studies that involve large displays, researchers were looking at broadly comparing one display technology to another. Hence setups were relatively un-

controlled and specific display characteristics poorly documented. For example, few of the papers measured factors such as field of view, resolution, brightness, or contrast of the displays used. Hence, while results were interesting, it is difficult to draw lower level conclusions that allow us to build a general theory around observed phenomena. Third, when researchers have explicitly studied the properties of large displays in an attempt to formulate design principles for their use, they have relied mainly on subjective ratings of a very varied set of media content, making comparison and generalization of results to productivity tasks difficult.

In my work, I take a bottom up approach, isolating and understanding display characteristics, so that we can inform a more general theory of cognition, especially as it relates to the effects of physical properties of information. This theory, while partially motivated specifically by considering physically large displays, should be general enough to include many classes of display devices and technologies.

Three important factors that differentiate large displays from smaller ones are the number of pixels or screen resolution, visual angle or field of view, and physical size. When they consider large displays, many people think immediately of having more screen space, or more pixels to place more information. This is true of multiple monitor systems, created by adding more displays to traditional single display systems. It is also true of high resolution displays, created specifically to increase screen space and display more information. Another factor of importance is field of view. Large displays are not often placed at a distance that is proportional to their increase in size over small displays. Due to space constraints, they are typically relatively closer and cast a larger retinal image, thus offering a wider field of view. While a large amount of work has been done in comparing fields of view, few researchers have isolated the effects of physical size and distance. Understanding the effects of physical size is one of the focuses of my work.

2.3.1 Number of Pixels

Anderson, Colvin, and Tobler (2003) studied 108 users working on single monitor and multi-monitor configurations. They found that users on multi-monitor setups outperformed users on single monitors on every performance and preference metric they tested. They concluded that adding pixels with multi-monitor setups are cost effective even if tasks they support comprised only about 20 percent of overall work done. Grudin (2001) explains how additional pixels provided by multiple monitors can be partitioned to take

advantage of focal and peripheral awareness as users work on various tasks. In his paper, he makes several high level observations regarding the use of multiple monitor systems and speculates on how we can design tools to optimally leverage these new systems.

2.3.2 Field of View

When considering field of view, it is important to define precisely what display characteristics are being referred to. There are two field of view angles that must be considered. The display field of view (DFOV) is the angle subtended from the eye to the left and right edges of the display screen. For a 16" wide display placed 24" from the user's eyes, the DFOV is approximately 37 degrees. This angle is limited by the physical display width, and can only be increased by replacing the display hardware or moving the user physically closer to the display. The DFOV can be decreased by using a window that is narrower than the screen width. The geometric field of view (GFOV) is the horizontal angle subtended from the virtual camera to the left and right sides of the viewing frustum. This angle is under control of the virtual environment designer. Most reported literature does not make the distinction between DFOV and GFOV. In most work, the term field of view (FOV) usually refers to geometric field of view. In my work, I use it mainly to refer to the display field of view. However, since I keep a 1:1 correspondence in all my work, it could just as easily refer to the geometric field of view.

It has recently been reported that it is harmful to deviate from a 1:1 ratio of GFOV and DFOV (Draper, Viirre, Furness, & Gawron, 2001). Large deviations can cause either magnification or miniaturization of items in the virtual world, possibly leading to discrepancies between studies as well as contributing reliably to simulator sickness. Barfield, Lim, and Rosenberg (1990) reported that performance was best under mid-sized GFOV conditions (45 or 60 degrees) and worst under extreme GFOV conditions (30 or 75 degrees), when they had participants judge azimuth and elevation under different conditions of field of view. They concluded that this was because the former GFOVs are closest to the DFOV and therefore result in the least amount of distortion.

There has been much evidence that restricting field of view leads to perceptual, visual, and motor decrements in various kinds of performance tasks (Alfano & Michel, 1990; Hagen, Jones, & Reed, 1978; Hosman & van Der Haart, 1981; Patrick, Cosgrove, Slavkovic, Rode, Verratti, & Chiselko, 2000; Piantanida, Boman, Larimer, Gille, & Reed, 1992; Sandor & Leger, 1991), though there is some debate about what field of view pa-

rameters are optimal in design for computing tasks. Alfano & Michel (1990) had users perform a series of hand-eye coordination tasks using goggles that restricted the field of view to 9, 14, 22 and 60 degrees. The 60 degree field of view condition yielded significantly better performance than the others, but all of the FOV restrictions were reported to cause disorientation in the subjects' depth and size judgments. Dolezal (1982) described the effects of restricting field of view to 12 degrees, including disorientation, dizziness during rapid head movements, difficulty in tracking objects, and difficulty forming a cognitive map of unfamiliar places. He observed that hand-eye coordination is impaired, and that there was greatly reduced ability to integrate visual information across successive views. Note that the inability to form a cognitive map of unfamiliar places coincides with the decrement in the overlap of visual information across successive views.

Examining cockpit displays, Kenyon & Kneller (1993) conducted two studies on the effects of different FOVs on the control of roll motion in cockpit displays. Response time delay and errors were found to decrease significantly with larger fields of view. However, most of the performance benefits were found with 40 or 80 degree FOVs, and there was little improvement with the full 120 degree FOV condition. In his work, Chambers (1982) concluded that the optimal field-of-view for flight applications was about 90 degrees on a virtual display. Increasing the amount of peripheral information by increasing the field of view up to 90 degrees reportedly allowed the user to construct an overlapping sequence of fixations in memory, which led to faster cognitive map construction.

In another study, Wells and Venturino (1990) reported that there was no effect of FOV on performance with only 3 targets to process in a display, but performance was significantly degraded by fields of view of 60 degrees or less when they increased the number of targets in the display to 9. In their study, users moved their heads less with the larger fields of view, since more of the targets were visible simultaneously on the display via eye movements.

In a series of studies, Czerwinski, Tan, and Robertson (2002) used a large widescreen display called Dsharp (see Figure 2.10) to replicate findings in the literature suggesting that wide fields of view offered by large displays helps users perform 3D navigation tasks more effectively. Interestingly, they found that while the wide displays benefited all users, it benefited females so much so that the gender gap that existed on traditional displays disappeared. In follow up work, Tan, Czerwinski, and Robertson (2003) showed



Figure 2.10. The Dsharp widescreen display has facilitated studies exploring the effects of wide fields of view in productivity as well as spatial tasks.

that the effect came from the presence of better peripheral optical flow cues, which females relied on more heavily than males.

In summary, it appears that wider FOVs offered by larger displays provide more spatial cues to users and are important aids for many spatial tasks, helping especially with cognitive map construction as the visual complexity of a display or the demands of a task increase.

2.3.3 Physical Size

Physical size is an important cue to sensory and judgment processes in humans. For example, both infants and adults have been shown to exhibit preferences for larger objects in presentations (Silvera, Josephs, & Giesler, 2002). Other research suggests that physical height in males is positively correlated with physical attractiveness, income, and occupational status (Jackson, 1992). In fact, the tallest candidate has won all but 4 of the 23 US presidential elections prior to 1992 (Newsweek, 1992). Wearing height and weight enhancing clothing and apparatus, including headdresses and shoulder pads, are instances of trying to manipulate this bias with size-inspired threat and power displays (Campbell, 1976). Additionally, Josephs, Giesler, & Silvera (1994) found that a wide variety of judgments were strongly influenced by non-diagnostic physical size information through the application of a ‘bigger is better’ rule. They found that artificially increasing the physical size of a pile by attaching each sheet of actual work to an empty cardboard box drastically increased estimates of productivity and progress. It is thus not a new idea that physical size of information greatly affects the way humans respond to it.

Despite the large amount of work done in comparing FOVs, few researchers have isolated the effects of physical size and distance on the sense of presence or on task performance. In a series of studies, Simmons (2001) showed that users performed better on productivity tasks using large 21" monitors as compared to smaller ones. Although this work showed benefits of larger displays, Simmons explored only a small range of display sizes, each viewed at different visual angles and with different resolution. Also, she described the presence of effects without attempting to explain them.

Dixon (1999) used the vertical-horizontal illusion, in which people overestimate the vertical but not horizontal extent of objects, to study the differences in perceived proportions between small and large displays. She found that the difference in magnitude for this illusion was influenced by the physical object or image size and not the perceived depth of the display. Her results further suggest that the proportions of larger objects appeared more compressed when presented on smaller displays, but that this effect could be compensated for by stretching the vertical proportions of images on a small display.

To examine the psychophysical effects of distance and size, Chapanis and Scarpa (1967) conducted experiments comparing the readability of physical dials at different distances. They used dials of different sizes and markings that were proportional to the viewing distance so as to keep visual angles constant. Perhaps surprisingly, they found that beyond 28" away, dials adjusted to subtend the same visual angle were read more easily at greater distances. The effects they found, however, were relatively small.

In a more recent study, Patrick, Cosgrove, Slavkovic, Rode, Verratti, and Chiselko (2000) examined various display technologies, with comparable visual angles, and their effects on the spatial information users acquired by navigating through a virtual environment. They found that while users performed significantly worse in forming cognitive maps and remembering the environment on a desktop monitor, they performed no differently using a head-tracked head-mounted display or a large projection display. They attributed part of this effect to a higher level of presence afforded by the physical size of the large projection display, which compensated for the immersion afforded by the head tracking.

Despite small pockets of results that the current body of literature offers, there seems to be a gap in work isolating the effects of display size and distance, given a constant vis-

ual angle, for performance on productivity tasks. Because of the emergence of large displays in the workplace and in consideration of everyday desktop computing tasks, I have examined the effects of holding visual angles constant and varying only physical size. I have examined subjective responses as well as cognitive and social reactions to information, and show how they may be used to construct display systems that make people more productive.

*“Form follows function - that has been misunderstood.
Form and function should be one, joined in a spiritual union.”*

Frank Lloyd Wright

“We do not see things as they are, we see them as we are.”

Anais Nin

Chapter 3

Large Displays Improve Performance on Spatial Tasks

3.1 Introduction

Large wall-sized displays are becoming prevalent. Although researchers have articulated qualitative benefits of group work on large displays, little work has been done to quantify the benefits for individual users. Furthermore, within the work aimed at quantifying benefits of large displays, little has been done to understand physical size as an important display characteristic that affects task performance.

In this chapter, I describe a series of experiments comparing the performance of users working on a large projected wall display to that of users working on a standard desktop monitor. Because I was interested in isolating the effects of physical size, I kept the visual angle subtended from the user to each of the two displays constant by adjusting the viewing distances appropriately (see Figure 3.1). I also held other factors such as resolution, refresh rate, color, brightness, contrast, and content as constant as possible across displays. Since the information content shown by each of the displays was equivalent, it would be reasonable to expect that there would be no difference in performance on one display or the other. However, I will show that this is not the case, and that physical size is indeed an important display characteristic that must be considered as we craft our display systems.




Figure 3.1. Basic experimental setup maintaining visual angles between the small and the large displays by adjusting the distance appropriately.

3.1.1 Exploiting User Perception to Support Cognition

In studying user reactions to physical properties of information, there are several stages of visual perception that we must recognize (see Figure 3.2). First, we assume that the physical world exists and that its existence is independent of the observer. Objects in the physical world are generally referred to as the distal or physical stimuli. Hence, physical size describes the actual size of an object in the environment. The physical world is governed by a well-defined set of physical laws. For example, objects reflect light in a predictable, though complex, manner. This world can also be described fairly completely with a homogeneous Euclidean geometry.

The observer, present in the physical world, views the physical stimuli when light reflected from objects in the environment stimulate receptors on the retinal surface of the eye. These impinging patterns of light are referred to as the proximal or retinal stimuli. Retinal size, then, describes the size of the image cast on the retina. It is also sometimes measured as the visual angle, or field of view, of an object or scene. The retinal size is dependant both on the physical size of the object as well as the distance from which the object is being viewed.

Perception refers to the conscious experience of the physical object or scene. Physical objects and scenes can be observed directly. Retinal stimulus patterns can be observed by projecting light from physical stimuli onto a projection plane that represents the retinal surface. Perception, however, is a process internal to the human mind and cannot be directly observed. Instead, we observe behavioral responses to the stimuli in order to indicate that certain perceptions are occurring. In my work, I am interested in understanding




Figure 3.2. Various stages of human perception – the distal object reflects light and forms an image on the retina, which is interpreted to form the perceived image.

and exploiting these behavioral responses by controlling the stimuli that shape visual perception.

The perceived image results from a complex relationship between physical and retinal cues, as well as cues internal to the observer, such as context and semantic knowledge. Unfortunately, most researchers have worked with the implicit assumption that the perceived image is dependant only upon the size and content of the retinal image. They have studied in detail how changing the retinal size by varying the visual angle or field of view affects perception and task performance. They have also developed many techniques to present information in a manner that is most easily processed by the human mind.

In my work, I recognize that perceptual space is shaped by more than just the retinal image. In fact, increasing the size of the display surface can fundamentally change the user perception and interaction with information. Even though a given image may have the same theoretical information content on a small or large display viewed at the same visual angle, it may elicit different cognitive and social reactions, causing responses that lead to different levels of productivity on different displays. In particular, I will explore through a series of experiments how varying the physical size of displays, while keeping factors such as visual angle and information content constant, affects perception and task performance.

3.1.2 Summary of Experiments

In Experiment 1, I show that physical size indeed affects task performance and should be further studied.

In Experiment 2, I describe how these effects may be due to display size automatically biasing users into adopting different strategies to perform tasks. In fact, I show that smaller displays seem to bias users into adopting exocentric strategies and large displays seem to immerse users more within virtual environments, biasing them into using more efficient egocentric strategies.

In Experiment 3, I add support to this explanation by testing performance on intrinsically exocentric tasks in which users do not benefit from using egocentric strategies. I show that large displays and the resulting egocentric strategies do not aid performance on these tasks.

In Experiment 4, I apply this understanding to a more ecologically valid task. I show that large displays help users encode perceptual movement information more effectively when navigating 3D virtual environments. Additionally, I show that the effects of large displays are independent of those caused by interactivity.

In Experiment 5, I generalize the results from Experiment 4 to a mental map formation and memory task, again showing the benefits of using large displays as well as the independence of large display effects from interactivity effects.

Finally in Experiment 6, I show that these effects, though slightly dampened, are robust even in a commercial off-the-shelf virtual environment.

3.2 General Experimental Setup

3.2.1 Equipment

I used two displays for each of the experiments, an Eiki Powerhouse One LCD projector and a standard-sized desktop monitor. In the first three experiments, I used an 18" Sony Trinitron E400 CRT monitor as the desktop monitor. In the other experiments, I replaced this with an 18" NEC MultiSync 1810X LCD monitor. All displays ran at a resolution of 1024×768 , updated at a rate of 60 Hz, and were calibrated to be of roughly equivalent brightness and contrast. I mounted the projector from the ceiling and projected onto a white wall. The image projected on the wall was 76" wide by 57" tall (see Figure 3.1). The image on the monitor was adjusted to be exactly 14" wide by 10.5" tall. I set the two displays up so that when either display was viewed from a specific spot in the room, the visual angle and hence the size of the retinal image, would be identical. I assumed a com-




Figure 3.3. Top view schematic of the experimental setup. I maintained constant visual angles by varying display size and distance accordingly.

portable viewing distance of 25" for the monitor. In order to get an image of identical perceived size, the projection was set up to be 136" away from the user. The center points of all displays were set to be at seated eye-height, set to be 48" above the ground.

Since the environmental context around each display could potentially affect users, I decided to keep the context as constant as possible by moving only the displays within the environment rather than having the user turn to face a different display with different environmental context. Hence, I carefully marked the position of the monitor so that it could be moved in and out as necessary.

I ran the first three experiments on a single 800 MHz Dell computer equipped with a dual-headed nVidia GeForce2 MX graphics card. I controlled the activation and deactivation of the displays using the Windows 2000 multiple monitor API so that only one display was active at any given time. For these experiments, the user provided input using an IBM USB numeric keypad with keys I had marked for the experiment (see Figure 3.4). I ran the latter three experiments on a 1.33 GHz computer with a GeForce4 MX graphics card. The virtual environments updated at 60 frames per second. I used a switchbox to send the graphics output to only one of the displays at any given time. The user provided input with the control stick and trigger button on a Radioshack 26-444 Joystick (see Figure 3.13).

3.2.2 Keeping Color, Brightness, Contrast Constant

I did several things to equate display characteristics such as color, brightness, and contrast across the various displays. Initially, I used a spectral radiometer and a colorimeter to measure the spectral distribution of the light coming off the displays as well as the tristimulus values of this distribution when various images were displayed. Unfortunately, as observed by MacIntyre and Cowan (1992), calibration done to an exact radiometric or colorimetric standard is both expensive and laborious. This is especially true of my setup, in which I was trying to calibrate different display technologies. Calibration is further complicated by human visual phenomena such as light, dark, chromatic, or transient adaptations (Milner & Goodale, 1996).

In my final calibrations, I took Tjan's (1996) view that a "human observer is always needed to carry out a color matching experiment." In fact, I assumed this to be the case for brightness and contrast as well. In order to calibrate the displays, I had groups of people view the two sets of displays. With questions such as "which screen do you think is brighter?" or "which screen has better contrast?" I was able to adjust the settings on the displays to be as constant as I could get them. I iterated with this process until users could not make these distinctions between the displays. It is also worth noting that the quality of the large projection display was probably poorer than that of the desktop monitor in all these regards. There is little reason to believe that the degraded quality would elicit any of the effects that we saw in the experiments.

3.2.3 Keeping Users' Heads Still

Another concern with the setup was that the visual angle calculations were only valid for a single point in the room. This meant that if users moved their heads from that point, the visual angles were no longer maintained between the two displays. This would cause complication in interpreting results. Even though the most controlled solution would have been to somehow fasten the user's head in place to prevent any movement, I decided against this because it would make the experiment both uncomfortable and unrealistic.

Instead, I marked the spot around which the user's eyes should have been centered by stretching fishing line from two stands, one on either side of the user. A mark in the center of the line indicated the exact spot in the room where the retinal images would be of identical size. For each user, I adjusted the chair so that they were seated comfortably with their eyes as close to the spot as possible and told them not to further move the

chair. I then removed the fishing line. In the rare case where users moved their heads or chair too much during the study, I readjusted their position before proceeding.

Pilot-test video showed that users hardly moved their heads after this initial adjustment during the study. In fact, most users never moved more than 2" to 3" in any direction. At various stages in this work, I also ran informal tests to validate experimental results when users' eyes were either a little too close or too far from the desired point in the room and saw similar effects to those observed in the experiments. Hence, I am fairly confident that the small head movements permitted within the setup did not directly account for the effects seen across the experiments.

3.3 Experiment 1: Physical Size Matters

Rapidly prototyping experiments when trying to design one is just as useful as rapidly prototyping interfaces when designing an interface. Experiment 1 was an exploration into the experimental design space. I chose a spatial and a textual task in order to determine whether display size had any performance effects on these general classes of tasks.

3.3.1 Participants

Twenty-four (12 female) college students, who were intermediate to experienced computer users, participated in the study. I screened users to be fluent in English and to have normal or corrected-to-normal eyesight. The average age of users was 25.4 (25.5 for males, 25.3 for females), ranging from 19 to 32 years of age. The experiment took about an hour and a half and users were paid for their participation.

3.3.2 Procedure

After users filled out a background survey, I gave them the numeric keypad and had them sit comfortably in the chair (see Figure 3.4). As previously described, I adjusted the



Figure 3.4. (left) Numeric keypad input device used in the first three experiments.
User working on the small (center) and large (right) displays.




Figure 3.5. Sample question from the Guilford-Zimmerman Spatial Orientation test.
The correct answer for this question is option 5.

height and position of their chair so that the center of their eyes was as close to the marked fishing line as possible. Once they were viewing the displays from the spot in the room that provided retinal images of identical size, I removed the fishing line. At this point, I instructed users not to further adjust the chair or move it around.

3.3.2.1 *Guilford-Zimmerman Spatial Orientation Task*

To evaluate the effects of display size on spatial performance, I utilized the Guilford-Zimmerman Spatial Orientation test (Guilford & Zimmerman, 1948). This test has been well validated and researchers have shown that results from this test correlate highly with wayfinding ability (Infield, 1991).

Each question in this test contained two pictures seen from the prow, or front, of a boat along with a multiple choice answer key (see Figure 3.5). The user was asked to imagine that each picture was taken with a camera fastened rigidly to the boat so that the camera bobbed up and down, slanted, and turned with the boat. First, the user looked at the top picture to see where the boat was initially heading. This heading is represented by the dot in the answer key. Next, the user looked at the bottom picture and determined the change in orientation of the boat. The line in each of the possible answers represents the new orientation of the boat relative to the previous heading. Finally, the user selected the answer with the number keys, confirmed the answer with the enter key, and proceeded to the next question. The full set of questions can be found in Appendix A.6.

In this experiment, I gave users the paper-based instructions that were provided with the standard Guilford-Zimmerman test. They then tried 3 practice questions. For these questions, the system provided users with immediate feedback explaining the correct an-

swers. After they had completed the practice questions, users performed the test on both the small and the large display, which I will refer to as the Display Size manipulation. The order of Display Size was counterbalanced across users. Users were not given feedback for the test questions. The 60 test questions were randomized and broken into two sets. Users had 5 minutes to answer 30 questions in each of the two conditions, and were told to work as quickly and accurately as possible. Users had a 30 second rest interval between each condition.

3.3.2.2 Reading Comprehension Task

I also tested subjects on a reading comprehension task in the two Display Size conditions. Based on the normalized average scores for the specific passages as well as pilot test data, I chose a suite of 7 passages from practice GRE tests (Educational Testing Service, 1994) that had relatively similar levels of difficulty. Each of these passages contained about 460 words, or 56 to 60 lines of text as laid out in paper-based GRE format. Each passage came with a set of 7 reading comprehension questions that the user answered after reading the passage.

I gave users verbal instructions on how to scroll through the passage and to answer questions. Then they performed the task with a practice passage on the large display with medium-sized text. I instructed them to work through the questions quickly but accurately. When they had finished the practice passage, they read the rest of the passages and answered questions in each of the 6 conditions, created by presenting text in a given Font Size (small font: 10 point vs. medium font: 14 point vs. large font: 18 point) on each of the Display Size conditions (small vs. large). Font Size and Display Size were counterbalanced separately. Again, users had a 30 second rest interval between passages.

3.3.2.3 Post-test Preference Questionnaire

After users completed the tests, they filled out a questionnaire indicating their preference for the conditions in each of the tasks. They were also encouraged to comment on their opinion of the displays.

3.3.3 Results

I present the results from this experiment in three parts. First I explore performance on the spatial orientation task, then the performance on the reading comprehension task, and finally I investigate preference measures collected at the end of the study.




Figure 3.6. Main effect of Display Size. Users performed significantly better on the Large Display than on the Small Display.

3.3.3.1 Spatial Orientation Task Performance

I analyzed data for the spatial orientation task at the summary level. The dependent variable was the percentage of correct responses (number correct / number attempted). Time differences between different Display Sizes were not significantly different and were therefore dropped from the final models. Levels of significance did not change either way. I analyzed the percentage of correct answers with a 2 (Display Size) \times 2 (Position) \times 2 (Gender) repeated measures analysis of variance (RM-ANOVA). I analyzed Gender and Position as between-subjects factors and Display Size as a within-subjects factor.

I found a significant main effect of Display Size ($F(1,20)=9.470$, $p=.006$) with the large display resulting in a higher percentage of correct responses than the small on average (55.4% vs. 43.8%, respectively; see Figure 3.6). I also observed a significant main effect of Gender ($F(1,20)=5.072$, $p=.035$), with males producing a higher percentage of correct responses than females on average (60.4% vs. 38.7%, respectively). None of the 2-way or 3-way interaction effects was significant.

I used percentage of correct answers as the dependent variable since it is the most straightforward and intuitive measure. Since this was a timed task, an alternate explanation for these findings may include a speed-accuracy tradeoff. However, a separate analysis confirmed there was no difference in time spent per question in the two conditions. I also examined the sum of correct responses, controlling for time, and found nearly identical results. Hence, the alternate speed-accuracy tradeoff explanation seems unlikely.

Overall, I found a significant improvement in the percent of correct responses on the spatial task for users working on the larger wall display. Keep in mind that while the absolute size of the image was larger, the perceived or retinal image size was kept nearly constant regardless of Display Size.

3.3.3.2 Reading Comprehension Performance

In the reading comprehension task, I again analyzed data at the summary level. I used the number of correct responses for each condition as the dependent variable. I performed an RM-ANOVA in which Position, Display Size, and Font Size were repeated and Time to complete the question was a covariate. I included all 2-way and 3-way interactions in the analysis. Because each user participated in multiple trials, within observations were not independent. I modeled User as a random effect.

Overall, performance in the reading comprehension task did not differ across the conditions. I found no difference between small ($M=3.86$) and large ($M=4.01$) Display Sizes ($F(1,106)=.367$, $p=.546$). Similarly, I found no difference between the small ($M=3.84$), medium ($M=3.96$) and large ($M=4.01$) Font Sizes ($F(2,106)=.176$, $p=.839$). The interaction between Display Size and Font Size was not significant, ($F(2,106)=1.159$, $p=.3178$).

I was unable to reject the null hypothesis that the displays were equal for performance on the reading comprehension task. Thus, while I did find differences on the spatial orientation task, I found no evidence to suggest that the performance on reading comprehension was different on either of the two Display Sizes, regardless of Font Size.

3.3.3.3 Preference Data

In addition to the performance data, I gathered preference data from users at the conclusion of the study. I asked questions on a 5-point Likert scale of 1=“Strongly prefer small display” to 5=“Strongly prefer large display”.

Given the performance difference I found on the spatial task, I was primarily interested in user preference for this task. Users significantly preferred the large display for both ‘Ease of Seeing,’ ($M=3.61$, $p=.019$) and ‘Overall Preference,’ ($M=3.50$, $p=.045$). They marginally preferred the large display for their ‘Confidence in the Rotation Task,’ ($M=3.43$, $p=.066$). Users showed no significant preference for display in the reading comprehension task.

3.3.4 Summary

In this experiment, I demonstrated the benefits of the larger display for performing a spatial orientation task, but found no evidence to suggest that reading comprehension was better in either display condition. The fact that I found differences in the spatial task but not the reading comprehension task led me to believe that there may be an interaction between the task and the display size.

While showing the presence of an effect, this experiment did not explain what caused the performance benefits when users worked on large displays. In order to further explore the reason behind this dramatic improvement in performance (approximately a 26% increase) I decided to run a second experiment to investigate the difference.

3.4 Experiment 2: Large Displays Bias Users into Egocentric Strategies

One explanation that accounts for performance differences in spatial orientation tasks is the choice of cognitive coordinate systems used to perform the task. This choice usually has implications on the particular strategy and hence the efficiency of performing the task. Just and Carpenter (1985) propose two strategies that might be used to perform the Guilford-Zimmerman test: an egocentric strategy and an exocentric one. Users performing the task egocentrically take a first-person view and imagine rotating their bodies within the environment. Users performing the task exocentrically take a third-person view and imagine objects rotating around each other in space. There is reasonable evidence in psychology research suggesting that egocentric strategies are more efficient for real world tasks (e.g. Carpenter & Proffitt, 2001). Hence,

Hypothesis 2a: Simple instructions and training prior to the test are sufficient to bias users into adopting either the egocentric strategy or the exocentric one when they perform the task.

Hypothesis 2b: The egocentric strategy is more efficient than the exocentric one for this spatial orientation task.

The instructions for the Guilford-Zimmerman test are carefully worded so as not to bias strategy choice one way or another. This allows users to either imagine themselves on the boat looking through the camera as the boat moves within the environment (egocentric),

or outside the environment as the boat rotates within it (exocentric). I believed that as users became more immersed in the task on the large display they were more likely to adopt the egocentric strategy. Since egocentric rotations have been shown to be quicker, this could explain the performance increase I observed on the large display. Thus,

Hypothesis 2c: With no explicit strategy provided, display size automatically biases users into adopting one or the other of the strategies. Small displays bias users into adopting an exocentric strategy, and large displays bias users into adopting an egocentric strategy.

3.4.1 Participants

Forty-two (18 female) college students, who did not participate in the first experiment, participated in this one. As before, I screened users to have normal or corrected-to-normal eyesight. The average age of users was 21.8 (21.7 for males, 22.2 for females), ranging from 18 to 35 years of age. The experiment took about an hour and users were paid for their participation.

3.4.2 Procedure

I used the same hardware setup as in the previous experiment. Recall that the instructions for the original Guilford-Zimmerman test were carefully crafted not to bias a user into any particular strategy. From this instruction set, I created two others, one that intentionally biased users into an egocentric strategy and another that biased users into an exocentric strategy (see Appendices A.3, A.4, and A.5). The egocentric instructions describe a scene in which users are asked to imagine themselves physically on the boat as it moves within the environment. The exocentric instructions describe the boat as a rigid prop mounted to the ground with the scene on a backdrop that is moving with respect to the boat. After balancing for Gender, each participant was randomly assigned to one of the three Instruction Types: Egocentric Instructions, Exocentric Instructions, or Original Guilford-Zimmerman Instructions. The overall procedure was the same as in Experiment 1, minus the reading test.

3.4.3 Results

I present results from this experiment in two parts. First I explore performance on the spatial orientation task, and then I examine the preference data.




Figure 3.7. Main effects of Strategy, with users performing significantly better with Egocentric Instructions than Exocentric ones. Also, results suggest that users with Unbiased Instructions perform with exocentric strategies when using the Small Display, and with egocentric strategies when using the Large Display.

3.4.3.1 Effects of Strategies on Task Performance

I modeled the data as I did for the spatial task in the previous experiment. I examined data at the summary level and used the percentage of correct responses (number correct / number attempted) as the dependant variable. I analyzed the percentage of correct responses with a 2 (Display Size) \times 3 (Instruction Type) \times 2 (Position) \times 2 (Gender) repeated measures analysis of variance (RM-ANOVA). I analyzed Instruction Type, Position, and Gender as between-subjects factors and Display Size as a within-subject factor.

Overall, I found a significant effect of Instruction Type ($F(2,37)=3.866$, $p=.030$; see Figure 3.7). Paired comparisons using the Bonferroni technique showed a significant difference between the egocentric and the exocentric instruction sets ($p=.01$), with users getting a higher percentage of questions correct with egocentric instructions than the exocentric ones (66.5% vs. 47.2%, respectively).

I conducted post-hoc tests to see if users who were explicitly instructed to use a given strategy performed any differently from users who implicitly chose a strategy due to the Display Size. I found no significant differences between users in the exocentric condition and the unbiased small display condition, which was assumed to elicit an exocentric strategy ($t(40)=.079$, $p=.9371$). Similarly, I found no significant differences between users in the egocentric condition and the unbiased large display condition, assumed to elicit an egocentric strategy ($t(40)=0.953$, $p=.3463$). I also conducted additional tests comparing performance on the small display in the exocentric condition to the small display in

the unbiased condition, as well as the large display in the egocentric condition to the large display in the unbiased condition. In both cases, there were no significant differences.

These results, seen in Figure 3.7, replicate findings from the previous study as well as provide support for my hypothesis that large displays provide a greater sense of presence and bias users into using egocentric strategies.

3.4.3.2 Preference Data

As in Experiment 1, I gathered preference data from the participants at the conclusion of the experiment. The merged preference data for all three Instruction Type conditions were not significantly in favor of the large display. I explored whether or not users in the different Instruction Types viewed the value of the displays differently.

I found in paired comparisons using the Bonferroni technique that users with the egocentric instructions and the unbiased instructions preferred the large display significantly more than users given the exocentric instructions for ‘Ease of Seeing’ ($p=.034$ and $p=.046$, respectively) and marginally significantly more in ‘Confidence in Rotation’ ($p=.064$ and $p=.077$, respectively). However, I did not see any significant differences in ‘Overall Preference’ across Instruction Types, suggesting that effects were probably not driven entirely by display characteristics and subjective preference. In general, these satisfaction ratings complement performance results nicely.

3.4.4 Summary

Results from this study show much more clearly that users performed better when they were provided with an egocentric strategy than when they used an exocentric one. Also, simple instructions and training were indeed sufficient to bias users into adopting one or the other of the strategies. In the absence of an explicit strategy, users seem to have chosen an exocentric one when working on the small display and the much more efficient egocentric one when working on the large display. Results from these first two experiments together suggest that, given a constant visual angle, the size of a display affects perception and performance in spatial orientation tasks.

In the next experiment, I provide additional support and insight into this explanation. If the explanation is correct, and the cause of the observed performance benefits is the implicit choice of an egocentric strategy, we would expect not to see benefits in tasks for which egocentric strategies do not help.

3.5 Experiment 3: Large Displays Do Not Aid Exocentric Tasks

While Guilford (1972) considered a single spatial orientation factor in his Structure of Intellect model, other researchers (e.g. Lohman, 1979) have identified three related spatial ability factors: spatial egocentrism, the ability of the observer to imagine their body in a different position so that they can figure out how a stimulus array will appear from another perspective; spatial relations, the ability to identify a certain objects when seen from different positions; and visualization, the ability to form a mental image of something that is not visible.

The Guilford-Zimmerman Spatial Orientation test used in the first two experiments allowed the user to use either spatial egocentrism or exocentric spatial relations strategies to perform the task. It was the choice of these strategies, biased either by prior instructions or by the size of the display that accounted for the observed performance differences. In this experiment, I picked tasks that did not seem like they would benefit from doing the task with a spatially egocentric strategy. Thus,

Hypothesis 3: Large displays bias users into using egocentric strategies and do not increase performance on ‘intrinsically exocentric’ tasks for which egocentric strategies are not useful.

3.5.1 Participants

Twenty-four (12 female) college students, who did not participate in the previous experiments, participated in this one. As before, I screened users to have normal or corrected-to-normal eyesight. The average age of users was 24.1 (24.4 for males, 23.8 for females), ranging from 18 to 44 years of age. The experiment took about an hour and users were paid for their participation.

3.5.2 Procedure

I used the same hardware setup as in the previous experiments. The tasks used in this experiment were selected because they are fairly abstract tasks and do not seem like they would benefit from having the user imagine their body within the problem space. The first two tasks, the Card test and the Cube test, are subtests S-1 and S-2 of the ETS Kit of Factor-Referenced Cognitive Tests (Ekstrom, French, Harman, & Dermen, 1976). The tests are inspired by Thurstone’s cards and cubes (Thurstone & Thurstone, 1941). The

third task, the Shepard-Metzler test (Shepard & Metzler, 1971), is a task commonly used to study mental rotation. I used a subset of questions from this test.

Before beginning the tasks, subjects filled out a background questionnaire and adjusted themselves in the chair so that their eyes were as close to the appropriate point in the room that ensured equivalent visual angles between displays. Subjects then did each of the Card Test, the Cube Test, and the Shepard-Metzler Test, in that order. This experiment was a within-subjects design, with each subject performing each of the tasks in both the Large Display and Small Display conditions in an order that was counterbalanced between subjects. Finally, they completed a preference survey.

3.5.2.1 Card Test

In each question of this test, the user saw two cards, each with the image of an irregular shape (see Figure 3.8). The two cards showed either the same shape or mirror images of the shape, rotated to different degrees. The user's task was to mentally rotate the cards in the plane and determine if they represented the same shape or if they were mirror images of each other. For detailed instructions, see Appendix A.7.

The original paper-based test presented a single base image to which eight other images were compared. Each section of the test was printed on a single page with 10 such rows of questions, for a total of 80 questions. In the computer-based version of this test, I showed each pair of cards one pair at a time, advancing to the next pair only when the user responded to the question. The left card in each pair corresponded to the base shape in the paper-based test. Users had 3 minutes to complete each set of 80 questions, seen in Appendix A.8.

3.5.2.2 Cube Test

In each question of this test, the user saw two cards, each with the drawing of a cube containing different characters in different orientations on each face (see Figure 3.8). Users were told that no character appeared on more than one face of a given cube. The user's task was to mentally rotate the cubes and determine if the drawings could have represented the same physical cube, or if they were definitely different cubes. For detailed instructions, see Appendix A.8.

Similar to the Card test, the paper-based test presented each set of 21 distinct pairs simultaneously on a single page. In the computer-based version, I showed each pair one

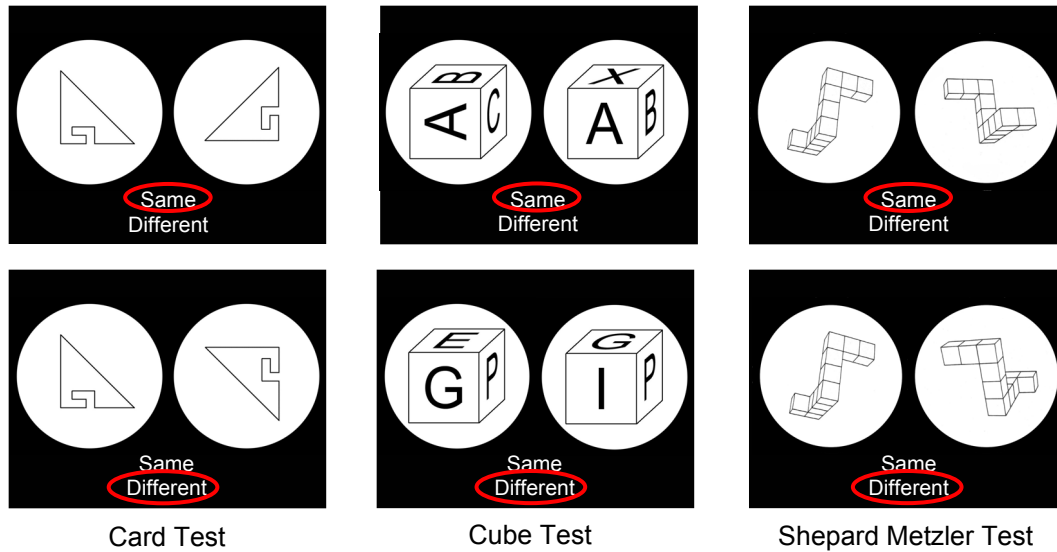


Figure 3.8. Exocentric tests that do not benefit from egocentric strategies, or user imagining their bodies within the problem space. In each, the user has to mentally rotate images to determine if they can be of the same object.

at a time, again advancing only when the user had provided an answer. Users had 3 minutes to complete each set of 21 questions, seen in Appendix A.10.

3.5.2.3 Shepard-Metzler Test

This test is similar to the Card test except that the mental rotation task is three-dimensional. Each question presents two drawings of objects in space (see Figure 3.8). Each object consists of 10 solid cubes attached face-to-face to form a rigid arm-like structure. Each is rotated to varying degrees. Users had to mentally rotate the objects in space in order to determine if they were the same object, or if they were different. For detailed instructions, see Appendix A.11. Once they indicated their answer using the keypad, the system advanced to the next question.

The original Shepard-Metzler stimuli of 70 line drawings consisted of 10 different objects in 7 positions of rotation about a vertical axis. These 7 positions permit the construction of at least two unique pairs at each angular difference in orientation from 0 to 180 degrees, in 20 degree increments. The full set of stimuli can be found in Appendix A.12. In this experiment, I created two equivalent subsets of the test, each with 60 questions: 6 objects \times 5 angles (20, 60, 100, 140 and 180 degrees) \times 2 answers (same or different object). I presented each pair to users one at a time. Users had no time limit for this task, but were reminded to perform the questions as quickly and accurately as possible.




Figure 3.9. Users performed no differently on any of the tasks whether using the Small or the Large Display. Egocentric strategies do not help on exocentric tasks.

3.5.3 Results

As before, I present results from this experiment in two parts. First I explore performance on the spatial orientation task, and then I examine the preference data.

3.5.3.1 Exocentric Task Performance

Since I did not expect any effects across experiments, I analyzed each of the tests independently. I modeled the data for each of the three tasks at the summary level. I analyzed the percentage of correct responses (number correct / number attempted) for each test with a 2 (Display Size) \times 2 (Position) \times 2 (Gender) repeated measures analysis of variance (RM-ANOVA). I saw similar results when I used the absolute number of correct answers as the dependent measure. I analyzed Gender and Position as between-subjects factors and Display Size as a within-subjects factor.

I saw no effects of Display Size in each of the three tests, with no significant difference in percentage of correct responses for the Card test ($F(1,19)=1.473$, $p=.240$), the Cube test ($F(1,19)=0.012$, $p=.914$), or the Shepard-Metzler test ($F(1,19)=0.5108$, $p=.475$). These results can be seen in Figure 3.9. Likewise, none of the other main effects or interactions was significant for this dependent measure.

When I compared the average time spent per question on each of the three tests, I found no significant interactions with the display manipulation. One point worth noting is that when I conducted an analysis at trial level, similar to that performed in the original Shepard and Metzler (1971) experiments, I found comparable results. I found significant




Figure 3.10. Results similar to Shepard Metzler (1971) experiments. The time taken for each questions is a linear relationship to the amount of rotation required.

effects ($F(1,2267)=32.704$, $p<.001$) suggesting that the larger the angle of mental rotation required to align the two objects, the longer it took users to decide whether the objects were the same or if they were different (see Figure 3.10). In fact, the relationship between angle of rotation and time spent on the question was a linear trend, as predicted.

3.5.3.2 Preference Data

Overall I found no significant differences in preference when users were asked to rate the two displays on a 5-point Likert scale of 1=“Strongly Disagree” to 5=“Strongly Agree”. The questions were ‘information on this display was easy to see’ ($M=4.33$ vs. $M=4.13$ for small vs. large display), ‘the task was easy to do on this display’ ($M=3.79$ vs. $M=3.70$ for small vs. large display), and ‘overall I liked this display’ ($M=4.13$ vs. 3.79 for small vs. large display). This corresponds well with performance data.

3.5.4 Summary

Even though there is evidence that the tests used in this experiment utilize similar cognitive abilities as the Guilford-Zimmerman task, namely spatial orientation and mental rotation, I asserted that users would not benefit from imagining their bodies within the problem space due to the abstract nature of stimuli. Results indeed showed that users did not experience the same benefits on these exocentric tasks that they did on the Guilford-Zimmerman task. In fact, users performed just as well when they worked on the small display as on the large display. This finding provides additional support to the explana-

tion that performance benefits were due to an increased sense of presence which biased users into egocentric strategies, strategies which were not useful for intrinsically exocentric tasks. It also implies that we must be very careful in applying the finding as large display benefits only apply to tasks which can be performed more effectively using egocentric strategies.

It was initially advantageous to use well validated and established psychology tests in order to understand the particular psychophysical phenomena in which I was interested. Although effects were easy to interpret, these tests had several shortcomings: (1) they were designed to isolate and study very controlled spatial abilities and did not take into account tasks in which compound abilities would be used; (2) the stimuli were often contrived two-dimensional, black and white images; and (3) they were static multiple choice tests that did not require the user to interact with the virtual environment.

In the following experiments, I extend the work by applying findings to more ecologically valid tasks. I incrementally increase the complexity of spatial abilities used in order to see if the current effects continue to be robust. I also use fairly rich dynamic three-dimensional virtual environments and incrementally increase the complexity of these environments by adding cues such as distinct landmarks and textures in order to see how the effects hold up in the presence of other cues. Finally, I test for the reliability of the large display effects when the user is actively interacting with the virtual environment. Interactivity could potentially also immerse the user within a virtual environment and cause them to perform better, hence negating some of the benefits afforded by large displays.

3.6 Experiment 4: Large Displays Improve Path Integration Performance

In this experiment, I chose a 3D navigation task to address the questions of external validity and real world usefulness of prior results. When navigating, users continually update mental representations of their position and orientation within the environment. This is called spatial updating. Two ways users can perform spatial updating are piloting, using external landmarks to get their absolute position within the environment, and path integration, sensing self-velocity and acceleration to derive their position relative to some starting point (Golledge, 1999). Path integration allows travelers to integrate isolated views of the environment into a cognitive map which may be used for subsequent pilot-

ing. Initial work (Chance, Gaunet, Beall, & Loomis, 1998; Klatzky, Loomis, Beall, Chance, & Golledge, 1998; Loomis, Klatzky, Golledge, Cicinelli, Pellegrino, & Fry, 2001) has suggested that successful path integration requires proprioceptive and vestibular cues, cues that provide physical awareness of our body's position with respect to itself or to the environment. However, recent studies (e.g. Kearns, Warren, Duchon, & Tarr, 2002) have demonstrated otherwise, showing path integration to be effective using only visual cues. Interestingly, many of these studies have presented the virtual environments on either physically large or wide field of view displays.

Riecke, van Veen, and Bühlhoff (2000) used a large half-cylindrical 180 degree wide projection screen and demonstrated that visual path integration, without associated proprioceptive or vestibular cues, was sufficient for elementary homing tasks. They claimed that additional peripheral cues provided by the display aided task performance. In other work, Péruch, May, & Wartenberg (1997) used a large video-projector screen and found that users navigated equally well in various field of view conditions, suggesting that task performance was independent of field of view. However, they did not explicitly discuss the influence that the physically large display had in their studies. My work contributes to this growing body of research, demonstrating that physical display size influences performance on these tasks.

Additionally, in this experiment, I further explore the egocentric vs. exocentric hypothesis that I have proposed as an explanation for effects. Interestingly, there have been two mental models suggested in connection with performing path integration, a traveler-centered model and an environment-centered model. These models relate directly to the proposed dichotomy of possible strategies, differentiating egocentric from exocentric representations. For a review of work in this area see Rieser (1989). If previous results generalize and large displays provide a greater sense of presence, biasing users into adopting egocentric strategies, I would expect performance to increase on this 3D navigation task. Thus,

Hypothesis 4a: Users perform better in the path integration task when using a physically large display due to the increased likelihood that they adopt egocentric strategies.

In separate work, some researchers have found that the acquisition of spatial knowledge is facilitated by active navigation control (e.g. Cutmore, Hine, Maberly, Langford, & Hawgood, 2000; Philbeck, Klatzky, Behrmann, Loomis, & Goodridge, 2001). These researchers claim that proprioceptive cues provided by the input devices as well as cognitive benefits of decision-making immerse users more within the virtual environments and aid in encoding mental representations of the environments. Others however, have reported opposite results, showing that active control hurts performance in various navigation tasks (e.g. Booth, Fisher, Page, Ware, & Widen, 2000). Flach (1990) argues that the different results could be due to the tradeoffs imposed by control of attention, kinds of information available, sensitivity to information, as well as activities involved.

I decided to explore both how level of interactivity in the virtual environment affects navigation by path integration, as well as how it interacts with effects caused by varying the physical size of displays. While prior literature provides evidence of active control helping in some situations and hurting in others, I expected users to perform better when they had interactive control using the joystick due to the additional cues afforded by the physical manipulation. Therefore,

Hypothesis 4b: Users perform better in the path integration task when they are interactively moving themselves through the virtual environment.

Finally, I expected the benefits of the large display to be robust against other factors that could potentially provide a similar heightened sense of presence. Specifically,

Hypothesis 4c: The effects induced by physical display size are independent of those induced by interactivity.

3.6.1 Participants

Sixteen (8 female) college students participated in the experiment. Users were intermediate to experienced computer users who played an average of less than an hour of 3D video games per week. All users had normal or corrected-to-normal eyesight. The average age of users was 23.3 (24.4 for males, 22.3 for females), ranging from 19 to 29 years of age. The experiment took about an hour and users were paid for their participation.




Figure 3.11. First person view of arena with the pole that users saw when performing the tasks.

3.6.2 Task

I used a triangle completion task to test how each of our manipulations affected path integration. In this task, I led users along two legs of a triangle and then had them find their way back to their starting position unaided. I picked this task because it is simple, well defined, and ecologically inspired. It is also commonly regarded as the most direct way of measuring path integration ability (Fujita, Klatzky, Loomis, & Golledge, 1993). I believe that these results extend to more complex navigation tasks.

To isolate effects, I created a virtual environment that provided optical flow and




Figure 3.12. Diagram of terms used in the triangle completion task. Black lines represent the actual triangle; gray lines represent user response.

depth cues necessary for path integration, but that did not contain distinct landmarks used for piloting (see Figure 3.11). The environment was a circular arena with two concentric circles of trees. The inner circle bounded the navigation area. It was 16 meters wide and contained ten 4 meter tall trees that were evenly spaced along the circle. The outer circle was 22 meters wide, and contained ten 5 meter tall trees that were darker in color than the trees in the inner circle. Users in pilot tests complained that the environment seemed static and unreal. To address this concern, I animated the trees to gently sway in the breeze. The ground had a uniformly speckled texture. The maximum speed a user could move was 2 meters per second. The maximum turning speed was 30 degrees per second.

Each trial in the test consisted of two phases, the encoding phase and the return-to-origin phase. In the encoding phase, I led users along the first two legs of a triangle (see Figure 3.12). For each leg, they saw a pole at the next vertex of the triangle (see Figure 3.11). Their task was to turn and move to each successive pole. Users could only turn when they were standing at a vertex. Additionally, they could only move forward and backward in straight lines, and only while they were facing the next vertex. This prevented users from straying off the defined path. Upon getting to the last vertex, users began the return-to-origin phase. In this phase, the poles disappeared and users had to use the joystick to turn and move to the origin, using only the mental map they had constructed and the visual cues provided by the environment. Again, users could only turn when they were at the vertex. However, since they could move forward and backward, they could return to the vertex to adjust their response angle if they felt that it was not correct. They pressed the trigger on the joystick to indicate their answer when they were done navigating.

3.6.3 Procedure

Prior to the start of the study, users performed the Map Planning (SS-3) subtest from the ETS Kit of Factor-Referenced Cognitive Tests (Ekstrom, French, Harman, & Dermen, 1976). This well-validated test is commonly used to evaluate spatial ability skills. The study setup was similar to the first three experiments and can be seen in Figure 3.13.

After reading through instructions, seen in Appendix A.13, users performed a set of practice trials before beginning the actual test. In these practice trials, users saw an overview map of the triangle before performing the task. After each trial, they received feedback on the overview map showing where they ended up relative to the origin. Each of




Figure 3.13. The joystick used (left); User working on the small display (center) and the large display (right).

the six practice trials used a unique triangle that did not match any of the test triangles. To prevent users from becoming reliant on maps, they were warned that they would not have these maps during the test.

The study was a 2 (Display Size: small vs. large) \times 2 (Interactivity: passive viewing vs. active joystick) within-subjects design. Users performed six trials in each condition, corresponding to six triangle configurations created by permuting three Angle 1 values (60, 90 and 120 degrees) and two Leg 2 lengths (3 and 5 meters). Leg 1 was always 5 meters long. These triangles can be seen in Figure 3.14. Each was centered in the arena.

In the passive viewing condition, users had no control of their movement in the encoding phase. Instead, they passively viewed themselves moving along the first two legs of the triangle. I used a slow-in slow-out animation with linear acceleration to move the user at the maximum speeds. In the active joystick condition, users used the joystick to navigate the first two legs. In both conditions, users had joystick control in the return-to-origin phase. I balanced the order of the Display Size and Interactivity manipulations separately and fully randomized the order in which I presented different triangles in each condition.




Figure 3.14. Six different triangles tested in the Experiment 4.

Dependent measures included: (a) the overall distance-to-origin error, the absolute straight line distance between the point to which the user navigated and the actual origin; (b) the angle-turned error, the signed difference between the correct angle (Angle 2) and the angle the user turned; and (c) the distance-moved error, the signed difference between the correct distance (Leg 3) and the distance the user moved. These error measures can be seen in Figure 3.12.

3.6.4 Results

3.6.4.1 Overall Task Performance

In my primary analysis, I examined the distance-to-origin error as the variable of interest. I used a mixed model analysis of variance (ANOVA) in which Display Size and Interactivity were repeated and Gender was treated as a between-subjects factor. I included all 2-way and 3-way interactions in the analysis. Because each participant performed each condition, observations within a pair were not independent and I modeled User as a random effect nested within Gender. I originally included two covariates in the model: Time spent in the encoding phase, and the Spatial Abilities score. However, I removed these from the final analyses because they were not significant. The estimates and significance levels of the main factors of interest did not change in any significant fashion and the overall model fit was improved.

I found a significant main effect of Display Size ($F(1,339)=11.24$, $p<.001$), with the large display resulting in users having shorter error distances than the small one (2.88 vs. 3.48 meters, respectively). I also observed a significant main effect of Interactivity




Figure 3.15. Main effects of Interactivity and Display Size. Note that users did better with Passive Viewing. There were no interactions between manipulations.

($F(1,339)=12.38$, $p<.001$), with trials in the passive viewing condition demonstrating shorter error distances than trials in the active control condition (2.87 vs. 3.49 meters, respectively). See these results in Figure 3.15. I saw no interaction between Display Size and Interactivity, suggesting that the manipulations were independent of one another.

Examination of the effect of Gender on performance did not reveal a significant difference between males and females (3.12 vs. 3.24 meters, respectively, $F(1,14)=.07$, $p=.79$). Prior literature suggests differential effects of gender on performance with different fields of view (Czerwinski, Tan, & Robertson, 2002). However, I controlled field of view to be constant across displays and saw no interaction between Gender and Display Size. This is consistent with the findings from my previous experiments. No remaining interactions were significant.

3.6.4.2 Systematic Component Errors

To test for systematic performance errors, I decomposed the aggregate distance-to-origin error and individually examined the distance-moved error and the angle-turned error. I used the same model as in the primary analysis, but replaced the dependent variable distance-to-origin error with the distance-moved and angle-turned errors.

I found a significant difference in Display Size for the distance-moved error ($F(1,339)=4.314$, $p=.03$). Users consistently underestimated the distance in both conditions (mean of 1.17 meter undershoot, overall). However, they underestimated significantly more in trials with the smaller display than the large (1.31 vs. 1.03 meter undershoots, respectively). The effect of Interactivity, while trending in the expected direction, was not significantly different for this measure (1.10 vs. 1.24 meter undershoot for passive viewing vs. active control, $p=.28$). While the mean result across all conditions demonstrated an underestimation of the angle (1.43 degree underturn, on average), I found no significant differences across the conditions for angle-turned error.

3.6.4.3 Effects of Triangle Shape

To examine effects of the different triangle configurations, I performed an additional analysis to explore whether the correct distance and correct angles affected performance in any systematic way. I performed a similar analysis as in the previous sections with correct distance (Leg 3) and correct angle (Angel 2) added as independent variables. I also

examined interactions to determine if the Display Size and Interactivity manipulations were more or less helpful depending on the difficulty of triangles.

I found, holding all other variables constant, that for every meter the correct distance increased, users accumulated an additional 0.635 meters in the distance-to-origin error ($F(1,354)=12.70$, $p<.001$). An examination of the interactions revealed that correct distance did not differentially affect performance across the various conditions. Correct angle had little effect on overall performance ($F(1,354)=1.47$, n.s.) and had no significant interactions.

In a similar fashion to the breakdown I performed with the systematic component errors, I also looked at the effect of correct distance on the distance-moved error as well as the effect of correct angle on the angle-turned error. I found that for each meter the correct distance increased, users underestimated the distance by an additional 0.465 meters ($F(1,358)=137.90$, $p<.001$). Similarly, I found that for each degree the correct angle increased, users further undershot the actual angle by an additional 0.501 degrees. I found no differential effects across conditions. These results are consistent with previous research showing that triangle shape significantly affects error rates (Kearns, Warren, Duchon, & Tarr, 2002).

3.6.5 Summary

These findings provide strong evidence that users perform 3D navigation tasks involving path integration more effectively on physically large displays than on smaller displays, even when the same environments were viewed at equivalent visual angles. In fact, in my simple triangle completion task, users performed about 17% better on the large display, supporting hypothesis 4a. Since more complex navigation tasks involving path integration can be decomposed into a series of triangle completions (Golledge, 1999), I could imagine the improvements cascading and leading to much greater overall benefits of using large displays. While there could be other ways to increase navigation performance, few alternatives provide as simple an extension to current tasks and methods as increasing the physical size of displays.

The effects I observed might be explained by the hypothesis that large displays provide users with a greater sense of presence within the virtual environment, biasing them into using more efficient egocentric strategies. One concern with this explanation is that

other mechanisms, such as interactivity, may affect task performance by evoking similar strategies. These mechanisms might then negate the effects provided by the large display. Results show that this is not the case, and that effects induced by large displays are independent of those induced by differing levels of interactivity, confirming hypothesis 4c. This means that designers can safely use different control mechanisms and continue to experience the benefits of their large display systems.

However, contrary to my initial hypotheses 4b, findings suggest that active joystick control is detrimental in the set of tasks I tested. I believe that this negative effect can be explained by the attention-cue tradeoff imposed by the new interaction mechanism and environment. The unfamiliar task of using a joystick to navigate the 3D virtual environment required a great deal of attention for my users, who indicated that they did not normally play 3D video games. This additional attention requirement probably impaired the creation of mental representations during the encoding phase in the main study. Because of the disparate reports of the effects of interactivity in various navigation tasks, I advise that researchers examine this manipulation carefully for their specific tasks and demographic before designing any interface and display system.

There could be several reasons that an egocentric strategy could cause a performance increase on the large display. In a follow-up investigation, I examined whether the errors seen in the main study were mainly cognitive errors or mechanical control errors.

3.7 Experiment 4b: Classifying Path Integration Errors

The implementation of the triangle completion task in Experiment 4 contained two fairly distinct subtasks: wayfinding, which included sensing the outbound path, forming a mental representation of the environment, and then computing the return path; and locomotion, or actually executing that path with motor movements to control the joystick. Since the errors observed could have been a result of either of these sub-processes, I ran a follow-up investigation to test the contribution that locomotion had on the error. Specifically, I wanted to know how well a user could use the joystick to turn a specified angle and move a specified distance. Thus,

Hypothesis 4d: Users were proficient with the joystick and the errors observed in the main experiment can be attributed mainly to wayfinding, or cognitive, errors.

3.7.1 Participants

Eight (4 female) college students, who had not participated in the main study, participated in the follow-up. I selected users to be of approximately the same demographic as before. The follow-up took about half an hour and users were paid for their participation.

3.7.2 Procedure

In this experiment, I simplified the triangle completion task to reduce the wayfinding component and test only how accurate users were in using the joystick to turn and move specified angles and distances.

Before each trial, I provided users with the angle and distance that they would have to turn and move. I told users the angle that they would have to turn (e.g. 60 degrees to the right). Unfortunately, since the virtual world contained no absolute unit of distance, telling a user to move 3 meters, for example, would not have been very useful. Hence, I specified the distance a user had to move by having them first travel a path of identical distance.

To reinforce these specifications, I showed users an overview map containing two legs of equal length connected at a single vertex. The user's task when placed in the virtual environment was to move straight ahead along the first leg, learning the distance they would have to travel. Following this, they had to turn the specified angle and move a distance equal to that of the first leg in order to reach the end-point of the second leg. They hit the trigger on the joystick to indicate when they were done navigating.

Before the study, users read the instructions, tried six practice trials in which they received feedback, and then performed the test. I tested angles and distances that represented the range performed in the return-to-origin phase of the main study. Using angles of 60, 90, 120, or 150 degrees, and distances of 3, 4, 5, 6, 7, or 8 meters, I created twenty-four test trials. Users performed these trials only in the small display \times active joystick condition. I expected that the largest locomotion errors would occur in this condition since it was the one in which users made the largest overall errors in the main study. This would serve as a good estimate of how much the locomotion errors were contributing to the overall error.

3.7.3 Results

I calculated 95% confidence intervals for our dependent measures. The distance-to-origin error had an interval from 0.31 to 0.39 meters, while the magnitudes of the distance-moved error had an interval from 0.18 to 0.22 meters, and the angle-turned error from 2.31 to 2.75 degrees. When compared to the mean magnitude of errors from this condition before, 3.78 meters, 1.71 meters, and 31.52 degrees, respectively, it is clear that locomotion errors account for a very small portion of the overall errors. This confirmed my hypothesis that wayfinding errors accounted for most of the errors seen in the main study.

3.7.4 Summary

I conducted this follow-up investigation to examine whether the errors seen in the main study were mainly cognitive wayfinding errors or mechanical control errors. I found that mechanical control errors accounted for only a very small portion of the total error, indicating that most of the error could be attributed to cognitive processes. This is consistent with assumptions in prior path integration literature, which attribute all errors to the encoding process (Fujita, Katzky, Loomis, & Golledge, 1993). This indicates that to increase performance with these measures, designers should spend more time optimizing cognitive cues, rather than control mechanisms.

In the following experiments I explore how we can utilize large displays to further optimize cognitive cues in tasks that require compound spatial skills. I further examine the Display Size \times Interactivity interaction, especially in light of a task in which the Interactivity manipulation aids performance. And finally, I move to the studying these effects within much richer and less controlled virtual environments.

3.8 Experiment 5: Large Displays Aid Map Formation and Memory

Results thus far show that information presented on physically large displays provides a greater level of immersion and allows users to perform certain tasks more effectively than on smaller desktop displays, even when information is viewed at equivalent visual angles. Tasks I have described so far include mental rotation (Guilford-Zimmerman task) and 3D navigation, specifically path integration (triangle completion homing task). In this experiment, I will extend these results to include a mental map formation and memory task. In this task, the user explores a virtual world in order to build a cognitive map of the environment. Using this cognitive map, the user then navigates to several specified targets

as quickly as they can. Users who build and remember better cognitive maps should be able to navigate to the targets with shorter distances and in less time.

There exists a vast body of work on general principles in 3D navigation. Thorndyke and Hayes-Roth (1982), as well as many others (e.g. Ruddle, Payne, & Jones, 1999; Siegel & White, 1975; Waller, Hunt, & Knapp, 1998; Witmer, Bailey, Knerr, & Parsons, 1996), have studied the differences in spatial knowledge acquired from maps and exploration. Darken and others have explored cognitive and design principles as they apply to large virtual worlds (Darken & Sibert, 1993, 1996). All this work recognizes that 3D navigation is a complex cognitive task requiring the use of a series of interrelated spatial abilities. I believe that benefits of large displays for simple spatial tasks extend to more complex ones, hence,

Hypothesis 5a: Users perform better in mental map formation and memory tasks when using physically large displays due to the increased likelihood that they adopt egocentric strategies.

Also, I have previously shown that effects of interactivity are independent from those caused by display size when active interaction is detrimental to the task. In this experiment, I hope to show that the effects are independent even when active interaction aids task performance. Thus,

Hypothesis 5b: Users perform better in the path integration task when they are interactively moving themselves through the virtual environment.

Hypothesis 5c: The effects induced by physical display size are independent of those induced by interactivity.

3.8.1 Participants

Sixteen (8 female) intermediate to experienced computer users from the Greater Puget Sound area participated in the experiment. I screened users to be non-gamers who played less than 3 hours of video games per week. I also screened users to be fluent in English and to have normal or corrected-to-normal eyesight. The average age of users was 36.0 (33.7 for males, 38.3 for females), ranging from 19 to 47 years of age. The experiment took about an hour and a half and users were given software gratuities for their participation.




Figure 3.16. (left) First person view of the world, including walls, target, and fence. (right) Map view of an example world. The user never saw this view.

3.8.2 Procedure

I created five different 3D virtual worlds using Touchdown Entertainment's Jupiter game development platform. Each of these worlds was a square room with edges 30 meters long. To prevent users from wandering outside the room, it was bounded by a fence (see Figure 3.16). Seven walls were randomly distributed throughout the environment. To ensure a well-distributed pattern of walls, I ensured that: (1) the average length of walls, approximately 7 meters, was comparable across the worlds, and (2) that each quadrant of the world had a roughly equivalent number and length of wall segments. I then distributed four red target cubes, one in each quadrant of the world. Each cube was uniquely marked and could be identified by the number of dots (one to four) found on each of its faces. Within this world, users had the same basic joystick controls as in the previous experiment. They moved at a maximum speed of 2.5 meters per second, and turned at a maximum rate of 30 degrees per second.

I used a mental map formation and memory task to test how each of the manipulations affected the way users performed in various 3D virtual environments. I broke the task into two phases: the learning phase, and the recall phase. In the learning phase, I gave users 4 minutes to explore the world and learn both the structure of the environment as well as the location of the various target cubes within the world. In the recall phase, I placed users in random locations within the world and had them move to specified targets as quickly as possible. These random locations were chosen such that the optimal path to the specified target was always 20 meters long. Users were asked to find each of the target cubes twice, for a total of eight trials per world. Detailed instructions can be found in Appendix A.14.

Note that the environments did not contain any distinct landmarks or textures (see Figure 3.16). The only way to remember the location of targets was to build a mental map using the structure of walls within the environment.

Dependent measures in this experiment included: (a) the distance moved from the start-point to the target in the recall phase; (b) the time required for the user to find each of the targets.

As with the previous experiment, this was a 2 (Display Size: small vs. large) \times 2 (Interactivity: passive viewing vs. active joystick) within-subjects design. In the active joystick condition, users utilized the joystick to move themselves through the environment as they explored it in the learning phase. In the passive viewing condition, users had no control of their movement in the learning phase. Instead, they viewed a movie of themselves moving through the environment. In pilot tests, I used the output from one user's active joystick condition as the stimulus for another's passive viewing condition. However, I found that most users moved themselves in somewhat unpredictable motions through the environment. This either caused an unreasonably high level of motion sickness in viewers, or was so jerky as to be ineffective in helping to learn the environment. Hence for the passive viewing condition I scripted a smooth path designed to explore the environment by moving in between every pair of walls at least twice in a systematic manner.

Prior to beginning the test, users completed a background questionnaire as well as a spatial ability test. I used the Paper Folding test (VZ-2) from the Kit of Factor-Referenced Cognitive Tests (Ekstrom, French, Harman, & Dermen, 1976). This test is a well-validated spatial orientation test that is commonly used to indicate general spatial ability. Users were then given detailed instructions and performed the task in the tutorial world. Following this, users performed the task in all four conditions, each in one of the four different environments. The conditions and the specific worlds were balanced across users. Finally, users filled out a preference questionnaire.

3.8.3 Results

Since the design of this experiment was identical to Experiment 4, I analyzed the data in a similar manner. I used a mixed model analysis of variance (ANOVA) in which Display Size and Interactivity were repeated and Gender was treated as a between-subjects factor,




Figure 3.17. Similar to Experiment 4, main effects of Interactivity and Display Size. In this experiment, users benefited from Active Control. There were again no interactions between manipulations.

including all 2-way and 3-way interactions in the analysis. Since this was a completely within-subjects design, observations were not independent and I modeled User as a random effect nested within Gender. I originally included two covariates in the model: distance moved in the learning phase, and the Spatial Abilities score. However, I removed these from the final analyses because they were not significant. Distance moved in the learning phase was not significantly different in any of the conditions, and Spatial Abilities did not interact with any of the manipulations. The estimates and significance levels of the main factors of interest did not change in any significant fashion and the overall model fit was improved.

In this experiment, I found a significant effect of Display Size ($F(1,487)=26.745$, $p<.001$), with the large display resulting in users moving shorter distances to find the targets than the small one (35.31 vs. 39.93 meters, respectively). I also observed a significant effect of Interactivity ($F(1,487)=14.219$, $p<.001$), with trials in the active condition demonstrating shorter distances moved than trials in the passive viewing condition (20.94 vs. 24.30 meters, respectively). Note that this is an effect opposite to that found with the path integration task in Experiment 4. Unlike results found for that task, interacting with the environment aided users with the map formation and memory task. As before, I saw no interaction between Display Size and Interactivity ($F(1,487)=0.909$, $p=.341$), suggesting that the manipulations were independent of one another. These results can be seen in Figure 3.17. Finally, I saw a main effect of Gender ($F(1,487)=9.119$, $p=.003$), with males

performing better than females (32.25 vs. 42.97, respectively). I saw no interactions between Gender and any of the manipulations.

The results for time required to find each of the targets matched findings using the distance moved metric. I found significant effects of Display Size ($F(1,487)=71.179$, $p<.001$), Interactivity ($F(1,487)=38.026$, $p<.001$), and Gender ($F(1,487)=5.259$, $p=.022$). There were no other significant effects or interactions. This is not surprising, as many subjects moved around at or close to the top speeds even when they did not immediately know their way around the environment.

3.8.4 Summary

This experiment adds further validity to my previous findings. In this experiment, I continued to see benefits of using large displays even with a fairly complex task requiring the use of numerous spatial skills. Like other tasks that benefit from the use of large displays, the map formation and memory task benefited from having users adopt an egocentric frame of reference while navigating.

In this experiment, unlike in Experiment 4, I found that active navigation control helped users learn and remember environments more effectively. In fact, they performed about 10% better when they controlled their movement than when they watched a video of themselves moving through the environment. Interestingly, and importantly, effects of interactivity were still independent of effects induced by display size.

One shortcoming of this experiment is the fact that the virtual environments used were still fairly sterile and controlled. They did not contain distinct landmarks or textures, which would be expected to exist in a more ecologically valid environment. I did this so that I could better understand the nature of the task and basic results before moving into a more complex environment in which other factors could contribute to effects observed. I conducted the next experiment to explore how robust these effects were in the presence of a multitude of additional cues found in more typical virtual environments.

3.9 Experiment 6: Ecological Validity of Results

This experiment extends previous results by testing the effects of display size in a much more ecologically valid environment. Thus,

Hypothesis 6: Even in an environment crafted with cues such as distinct landmarks and rich textures to be realistic and memorable, users perform better in mental map formation and memory tasks when using physically large displays due to the increased likelihood that they adopt egocentric strategies.

3.9.1 Participants

Sixteen (8 female) college students who were intermediate to experienced computer users participated in the study. I screened users to be non-gamers who played less than 3 hours of video games per week. I also screened users to be fluent in English and to have normal or corrected-to-normal eyesight. The average age of users was 23.9 (24.6 for males, 23.1 for females), ranging from 18 to 31 years of age. The experiment took about an hour and a half and users were paid for their participation.

3.9.2 Procedure

I used the same mental map formation and memory task as in Experiment 5. In the learning phase, users were allowed to explore a virtual environment for 4 minutes. In the recall phase, they had to locate specific targets from random locations within the world. These locations were randomly distributed to be distances of 50 meters away from the targets.

I used an off-the-shelf copy of Unreal Tournament 2003 (Epic Games) for this experiment. Unreal Tournament is a first-person shooter like Doom or Quake (ID Software), and can be considered to utilize a state-of-the-art rendering engine and virtual environments. In fact, virtual environments in this game are specifically crafted to be realistic, or immersive, and memorable for players (see Figure 3.18). Unreal Tournament



Figure 3.18. First person view of the world, which contains distinct landmarks and rich textures in Experiment 6. The target is the red flag.

comes with a game-mode called ‘capture the flag.’ Each of the worlds in this game has two team flags. In order to score, one team must touch the enemy flag and return it to their home base. I used these flags as targets to find within the environment.

Unreal Tournament comes with development tools for editing maps as well as for scripting simple behaviors within the worlds. I made two modifications to the game in order to run the experiment. First, I instrumented the game so that I could log the dependent measures: (a) the distance moved from the start-point to the target in the recall phase; (b) the time required for the user to find each of the targets. Second, I had initially left several computer enemies in the game to serve as further distraction while the user performed the mental map formation and memory task. However, in pilot tests, users got so carried away chasing and shooting enemies that they forgot all about their main task. As such, I removed all enemy characters as well as weapon pickups from the worlds in the actual tests.

I chose worlds from the standard set of worlds that ship with the game as well as from upgrade packages created by gamers and distributed on various websites. Through pilot tests, I selected five of these worlds from an initial pool of twelve, such that I had one small tutorial world, two Easy Worlds, and two Difficult Worlds. The Easy Worlds both covered about 1000 square meters and the Difficult Worlds covered a little more than twice that amount of space. Additionally, the Difficult Worlds were much harder to learn and navigate due to the complexity of structures and cues within the environment. For example, one such world had a maze of underground caverns and tunnels to navigate. Pilot tests suggested that each pair of worlds, the Easy Worlds and the Difficult Worlds, were of roughly similar difficulty within our task.

I eliminated the Interactivity manipulation from the previous experiment, and hence this experiment was a 2 Display Size (small vs. large) \times 2 Difficulty (easy vs. hard) within-subjects design. The orders of Display Size and Difficulty were independently balanced between users.

3.9.3 Results

I used the same analysis model as in the previous two experiments, replacing the Interactivity manipulation with the Difficulty one in this experiment. While I saw dampened effect sizes from the previous experiment, possibly due to the dilution caused by the addi-




Figure 3.19. Main effect of Display Size. Users performed significantly better when working on the Large Display than the Small one.

tion of cues within the environments, I observed similar findings. I found a significant effect of Display Size ($F(1,604)=11.900$, $p<.001$), with the large display resulting in users detouring by shorter distances to find the targets than the small one (14.71 vs. 16.35 meters, respectively). I also observed a significant effect of Difficulty ($F(3,604)=108.996$, $p<.001$), with trials in the easy worlds demonstrating shorter distances moved than trials in the difficult worlds. I saw no interaction between Display Size and Difficulty ($F(3,604)=4.041$, $p=.007$), suggesting that the manipulations were independent of one another. Finally, I saw a main effect of Gender ($F(1,604)=6.699$, $p=.010$), with males performing better than females (12.94 vs. 18.14 meters, respectively). I saw no interactions between Gender and any of the manipulations.

Again, the results for time required to find each of the targets matched findings using the distance moved metric, with a significant effect of Display Size ($F(1,604)=4.281$, $p=.039$), Difficulty ($F(3,604)=294.510$, $p<.001$), and Gender ($F(1,604)=7.319$, $p=.007$), but no other main effects or interactions.

3.9.4 Summary

This experiment shows that benefits of large displays are independent of cues that may be used in real-world virtual environments to increase immersion and memorability, such as distinct landmarks and rich textures. This is an important property if we are to apply the summary of results to useful real-world tasks, such as training and simulation, or games and entertainment. Also, it implies that we can continue to exploit the benefits of large displays even in the presence of other techniques that induce performance increases.

3.10 General Discussion of Experimental Results

The series of experiments described in this chapter demonstrate that physical display size is an important factor to consider when designing display systems. Results suggest that physically large displays, even at identical visual angles as small displays, immerse users and bias them into adopting egocentric strategies. These strategies increase performance on spatial tasks such as 3D navigation as well as mental map formation and memory, which can be represented using egocentric coordinate systems. Furthermore, the effects caused by physically large displays seem to be independent of other factors that may induce immersion or increase performance. For example, even though interactivity and mental aids such as distinct landmarks and rich textures within virtual worlds increase task performance on the tasks tested, they did not affect the benefits that large displays offer to users.

In fact, with very little effort on the part of the designer, the system builder, or the user, large displays offer the potential to improve performance on a fairly broad range of tasks. Also, because effects are independent of other aids tested, large displays continue to offer improvements even in the presence of other performance aids.

It could be argued that the magnitude of effects was not amazingly large, and that 10% to 26% increases are not enough to warrant the additional cost and physical space that large displays require. However, given that the theoretical information content shown on the small and the large displays were the same, and hence that the retinal images created when viewing one display or the other was the same, it is interesting that these results exist at all.

Furthermore, it should be noted that performance gains even of this magnitude could be important in the domains for which I think these results are most useful, namely games and entertainment, as well as training and simulation. Games and entertainment is an already large market that continues to grow, and that could benefit significantly from even a small portion of the demographic preferring and upgrading to large displays. In training and simulation, any small increase in performance could potentially lead to fairly large implications. For example, imagine firefighters who could navigate to targets 10% quicker or could better find alternate routes when they become obstructed because they trained on large displays.

The behavioral effect and choice of different strategy depending on the physical size of the display is perhaps more interesting than the raw magnitude of performance increases. The magnitude of the effect is heavily dependent both on the particular task as well as the surrounding context in which the task is performed. However, the behavioral effect can be attributed to a much more fundamental cognitive mechanism, which may form an important component of the way we perceive and interact with the world around us. In fact, it is interesting how robust the results were, both to the types of tasks tested, but also to the demographic for which this applied. Because the experiments were performed both with college students from Carnegie Mellon University as well as with a wide range of people recruited from the general population in the Greater Puget Sound area, I can say with relatively high degree of confidence that the results are representative of a large portion of the population.

In fact, informal observations across the demographic yielded other interesting design considerations. For example, people with bifocals usually preferred reading and performing the tasks on the large display. This was because they were much more comfortable working on surfaces that were further away. Depending on demographic, users would compare the large display to a movie screen or a classroom board, but most users indicated that they were more engrossed by the large display. Unfortunately, it was difficult to get users to articulate the level of immersion or the actual strategy they used to perform tasks. In fact, in pilot experiments as well as in the actual experiments, I tried various methods including multiple choice and ranking questions, magnitude questions using Likert scales, subjective open-answer questions, and informal interviews. While preference ratings generally matched performance results, none of these methods was effective in deriving definitive responses or insights regarding strategy used. Instead, I had to resort to carefully designing the experiments such that I ended up with a series of performance results suggesting that the strategy hypothesis is the most likely explanation for the effects observed.

As a final note, although I observed effects of gender and spatial ability across many of these tasks, I did not pursue these further. These effects have been fairly well documented in the literature and were not the focus of my experiments. While Czerwinski, Tan, and Robertson (2002) suggest that females benefit significantly more than males in 3D navigation tasks using displays with wide fields of view, I saw no such effect for Display Size in our studies. I found no interactions between any of the manipulations and

these factors, indicating that nothing surprising was happening with these effects. I have found no evidence suggesting that physical display size aids any part of the demographic more or less so than any other group.

3.11 Future Work

Although I did not intentionally calibrate the absolute size of the images in any of the experiments, images shown on the large display were close to being life-sized. This might be an interesting point in the size-performance curve as it could represent the optimal size at which users will be immersed. However, more work is required to determine the shape of the curve as one increases display size from a traditional desktop display to a large wall-sized display and then beyond. It would be interesting to see if the strategy change is an abrupt shift that happens when a certain size is achieved or if it is more continuous across a series of sizes. Also, it would be interesting to see what happens when images get larger than life. This would allow us to gain a deeper understanding of display size and how it relates to immersion and presence.

Another potentially interesting realm of study has to do with the factors that best allow us to perceive physical size. There are numerous factors such as optical accommodation and convergence, stereo vision, parallax, and environmental context, but we do not have a clear understanding of how each of these contributes to the effects and how they interact with each other. I believe that this could form an entire research agenda in and of itself, as this would add not only to our understanding of large displays, but of the human visual and perceptual systems in general.

Building upon an understanding of what it is that allows us to perceive size, I believe that it is also important for us to completely understand what it is about that size that causes us to become more immersed and to adopt different strategies when performing spatial tasks. For example, if these results could be partially attributed to novelty of the large displays, then results would be a little less useful theoretically, but it would be very interesting to find out that they were due to certain fundamental biases in our neural circuitry. Again, this remains future work.

Before we fully understand the design principles derived from these experiments within real world scenarios, there are a few other areas to consider. For example, we must look beyond behavioral responses and performance results, and understand how large

displays affect other things such as simulator sickness. In my experiments, I saw no indication that large displays would cause any more or less illness, but I cannot draw any conclusions because almost no one got sick in the experiments.

Also, we must fully examine the interaction of other display characteristics, such as field of view or resolution. Maintaining constant visual angles was merely a means to isolate and study physical display size as an interesting characteristic. I do not propose that large displays should be intentionally used at equivalent visual angles to their smaller counterparts. Instead, we should clearly understand the interactions between display size and these other factors so that we can design display systems that make optimal use of large displays.

Finally, I believe that we must explore training transfer, both between different types of displays, but also from the virtual world to the real one. In all my experiments, I did not see any ordering effects, suggesting that the strategy change was a rather ephemeral one, changing quickly depending on which display the user was currently using. However, given the length of my tasks, I would be hesitant to draw conclusions about any longer term training transfer effects, and this will have to be studied in more detail.

“This [WinCuts] is Tivo on steroids. Tivo extends television and allows me to watch the shows I want, when I want. WinCuts extends computing and allows me to work with the content I want, where I want.”

Randy Pausch

Chapter 4

WinCuts – A Tool that Facilitates Collaboration on Large Displays

4.1 Initial Motivation

Many researchers have used large displays to support various kinds of collaborative work. Within this work, several researchers have explored the utility of creating workspaces in which users, each with a personal computing device such as a desktop machine or laptop, can visually share information on a large display. Much of the work done in this area has been supported by fairly large hardware and software infrastructures. In my work, I set out to build a tool that would allow co-located users to easily collaborate on large shared displays in a fluid manner and with minimal infrastructure requirements. This tool would allow users to easily share information and windows between multiple machines (see Figure 4.1) so that they can simultaneously place useful information on a large display and make it visible to others. Ideally, everyone would be able to interact with that information in a simple manner.

There are two problems that arise in trying to do this. First, sharing windows between machines is not well supported in most current computing environments. Second, screen space is scarce. In fact, even if users are able to place windows on the large display, they quickly run out of space to place and view new ones simultaneously. I assert that this is one instantiation of a much larger information management problem that exists even on standard desktop systems.



Figure 4.1. WinCuts was initially motivated by the desire to provide a tool that would allow multiple users to share information on a large display.

An increasing number of everyday computing tasks require that users coordinate and operate on information from multiple sources. Each source of information is typically contained within a window, the fundamental unit at which users can easily manipulate chunks of information. Oftentimes, users benefit from simultaneously viewing relevant information that exists within different windows. Additionally, the spatial layout of this information may be crucial to effective task performance as it helps users not only to establish spatial relationships but also to visually compare contents.

Unfortunately, even with the emergence of large high-resolution displays, there is usually not enough screen space to view all windows simultaneously. Even when this is possible, having entire windows visible can introduce so much extraneous space in between relevant information that spatial location becomes less helpful. Users can adjust windows so that they contain only the relevant information, and then lay them out. However, this is an extremely tedious task as it involves multiple iterations of resizing windows and scrolling content, especially in applications such as web browsers that recalculate the layout of information based on the size of the window.

My work in this area has focused on allowing users to easily specify and organize regions of information contained within multiple windows so that they can make more effective use of screen space to perform their tasks more efficiently. In this chapter, I describe a novel interaction technique that allows users to replicate arbitrary regions of existing windows into independent windows called WinCuts. Each WinCut provides continuous visual updates of the region as well as input redirection mechanisms that allow users to interact with content. I further describe how the ability to share WinCuts

across devices extends the utility of this system to multiple users engaged in co-located collaboration as well as to single users working with multiple devices. I provide a non-exhaustive classification of the set of tasks that I believe might benefit from the use of WinCuts. Next, I discuss high-level implementation details for my current prototype. Finally, I discuss future directions that I will pursue with this work.

4.2 Background

Many researchers have identified and studied the specific potential of media spaces for collaboration. For example, Hollan and Stornetta (1992) propose a framework for research aimed at making distance collaboration as productive as face-to-face collaboration. Gaver (1992) discusses the affordances offered by media spaces for collaboration, comparing and contrasting their properties with those of the physical environment and everyday media. Ackerman (2000) describes the social-technical gap that divides what we must support socially and what we can currently support technically.

To realize some of this potential, Stewart, Bederson, and Druin (1999) present the Single Display Groupware model for supporting collaborative work between physically co-located users. They define single display groupware to be computer programs that enable co-located users to collaborate via a shared computer and a single shared display as well as simultaneous use of multiple input devices. For a review of the state of single display groupware research, see Shoemaker (2001). In this paper, he hints at the questions of how physical instantiations of technology affect the acceptance by users. Interestingly, he also mentions the scarcity of screen real-estate when multiple users work on a single display as an important open research question.

Various classes of techniques have been used for alleviating the problem of scarce screen real-estate. The first class is ‘distorted views’ of the virtual space. One of the oldest techniques for distorting space is using symbolic representations of larger objects, such as icons (e.g. Goldberg, 1984; Smith, Irby, Kimball, Verplank, & Harslem, 1982). Researchers have also explored visual distortions of information, forcing all objects to fit on the screen by scaling objects based on their intrinsic importance or by the user’s focus of attention (e.g. Furnas, 1986; Spence & Apperly, 1982). This work has led to numerous visualization techniques meant to take full advantage of scarce screen space.

The second class of techniques aimed at solving the screen space problem is ‘large virtual workspaces’. In this work, researchers arrange objects on a single large virtual workspace much larger than the screen. The screen is then treated as a movable viewport into that space. Sketchpad (Sutherland, 1963) was one of the earliest graphical programs to use this technique. More recently, the Pad and Pad++ systems (Perlin & Fox, 1993; Bederson & Hollan, 1995) provide the user with an infinite two-dimensional information plane. Users manipulate ‘portals’ that allow them to navigate and peer into the desktop. One interesting property of this system is that the information plane can be shared among many users, potentially providing interesting mechanisms for sharing information and collaborating.

The third class extends upon this and incorporates multiple virtual workspaces that the user can manage and view, usually one at a time. One early instantiation of such a system was Rooms (Henderson & Card, 1986), which divides the virtual workspace into a series of ‘rooms’ with transitions between them. Principles derived from this system have been widely implemented in a myriad of virtual desktops today.

Within all the techniques for managing scarce screen space, there exists the need to easily manage information and individual windows. To address this need, many researchers have explored window management systems, which provide different mechanisms for users to arrange multiple windows on the screen (for history and review, see Myers, 1988). Within this body of work, many researchers have compared the cost-benefit trade-offs that such systems impose. For example, Bly and Rosenberg (1986) compared tiled window systems, which automatically determine the size and location of all windows such that each window is always completely visible, to overlapping window systems, which allow the user to control size, location, as well as the overlap or visibility of windows. They found that tiled systems were optimal when the system picked arrangements that conformed to the relevant contents of windows. However when they did not, overlapping systems were far superior, even though the user had to exert additional effort to explicitly manage windows. These results suggest that the ideal system is one which requires users to exert the least amount of effort but ensures that information is laid out in a manner that best supports the task at hand. Wickens and Carswell (1995) further augmented these findings with their proximity compatibility principle, which holds that the more two information sources are used within the same task, the closer they should be displayed on the screen.

More recent work such as the Adaptive Window Manager (Stille, Minocha, & Ernst, 1997), Elastic Windows (Kandogan & Schneiderman, 1997), and Hutchings' operations for display space management (Hutchings & Stasko, 2002) has explored different mechanisms for more efficient windows management. Unfortunately, since these schemes continue to treat windows as the fundamental unit of information, it is still difficult to lay out smaller chunks of relevant information contained within multiple windows.

In separate work, researchers have explored the benefits of viewing and operating on information and applications across multiple devices. This is useful both for individual users working on multiple devices, as well as for groups of users, each with personal devices, working together. Like much of the work on window management systems, most of this work allowed users to share entire screens or regions of the screen (Richardson, Quentin, Wood, & Hopper, 1998) or, more recently, entire windows and applications. For example, Gutekunst, Bauer, Caronni, Hasan, and Plattner (1995) present a system that allows users to simultaneously view and interact with a single application on multiple workstations. For more detailed review of such systems, see Li and Li (2002). There also exist numerous open source and commercial applications with provide this functionality (e.g. VNC, Xwin, Remote Desktop, Maxivista). Again, these allow users to share entire windows, screens, or regions of screens between different machines. To my knowledge, WinCuts is the first piece of work that allows users to explicitly manage and share smaller regions of windows.

4.3 WinCuts Interaction Technique

4.3.1 Basic Technique

To create a new WinCut, users hold down a keyboard modifier combination, control-“accent grave”, which brings up a semi-transparent tint over the entire desktop. They then click and drag the mouse over a region of a window to specify a rectangular region of interest (ROI). They can redefine this region as many times as they like. When they are satisfied with the ROI, they release the keyboard keys. The tint disappears and a new WinCut appears on top of the source window, slightly offset from the location of the ROI. The source window is unaffected. The WinCut is differentiated from regular windows by a green dotted line around the content region of the WinCut (see Figure 4.2). Users may make as many WinCuts as they wish, either from a single source window, or from multiple windows.

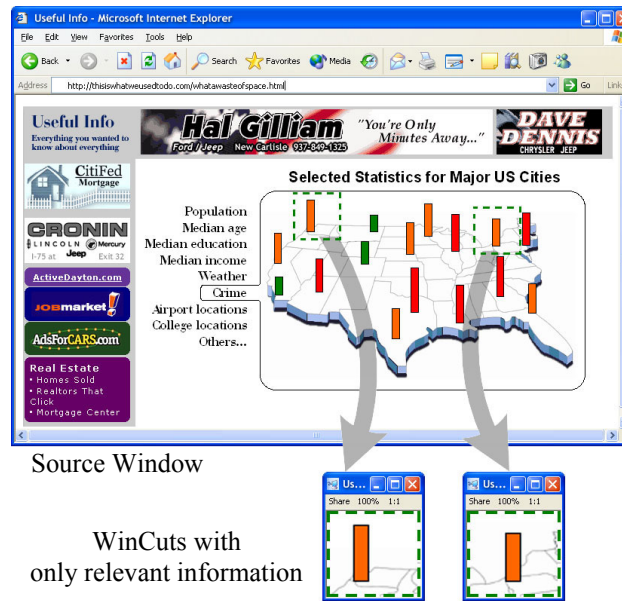


Figure 4.2. User makes two WinCuts to compare statistics between two cities. Each WinCut is an independent live region with which users can interact.

Each WinCut is a separate window and can be managed much like a regular window. It shows up in the Windows taskbar and can be minimized, restored, moved, and closed. Unlike other windows, however, maximizing or resizing a WinCut preserves the relevant information that is shown and instead rescales the content within the WinCut. This allows the user to make the information fill as little or as much space as they would like. For convenience, I have provided menu functions that allow a user to return the content to its original size or to constrain its aspect ratio.

WinCuts contain live representations of the content that appears within the ROI on the source window. In other words, the user can not only view updating content from the source window through the WinCut, but can also directly interact with content in it, just as they would in the original window. Since the content in each WinCut is tethered to the content of the source window and not a region of the screen, users can move and even hide the source window without affecting the WinCut.

4.3.2 Sharing WinCuts Across Machines

I have augmented this basic interaction with the ability to share WinCuts across multiple machines. Running the WinCuts system allows a user to send and receive WinCuts from other machines running the system. After creating a WinCut, a user can click on the “Share” button in the menu bar of a WinCut (see Figure 4.3). A dialog box pops up for




Figure 4.3. There are three buttons in the menubar of each WinCut: Share allows users to share WinCuts between machines, 100% returns the WinCut to original size, 1:1 returns it to original aspect ratio.

the user to specify the machine with which the WinCut should be shared. Doing this causes the current WinCut to appear on the destination machine. Once this is done, the user can interact with any local window and have the relevant WinCut update on the destination machine. However, aside from minimizing, restoring, resizing, moving, and closing the WinCut, users cannot currently interact directly with content on the remote WinCut. I do not perform input redirection on remote WinCuts because of the complication introduced when multiple input streams collide. A red dotted line around the content indicates that the WinCut is read-only on the remote machine. This system works with multiple machines simultaneously sending and receiving WinCuts.

In order to provide a simple mechanism to manage WinCuts on remote machines, I couple WinCuts with a separate program called Visitor, developed at Microsoft Research. Visitor redirects the input stream over the network so that a user can use a mouse and keyboard connected to one computer to control the input on another computer. This is similar to widely available utilities such as ShareKMC or Synergy. When running Visitor, a user simply moves the cursor off an edge of the local screen to take control of the cursor on a remote screen. In this way, users can easily use the mouse connected to the local computer to manage WinCuts on the remote screen. Multiple users can take turns doing this.

4.4 General Usage Scenarios

4.4.1 Single Machine Tasks

One motivation for creating WinCuts was to provide users with tools necessary for effective spatial organization of information with limited screen space. At its core, WinCuts provides a lightweight mechanism for a user to specify relevant regions of information contained in various windows and then use standard windows management techniques to organize and lay these out. The most fundamental use of WinCuts, then, is to specify re-

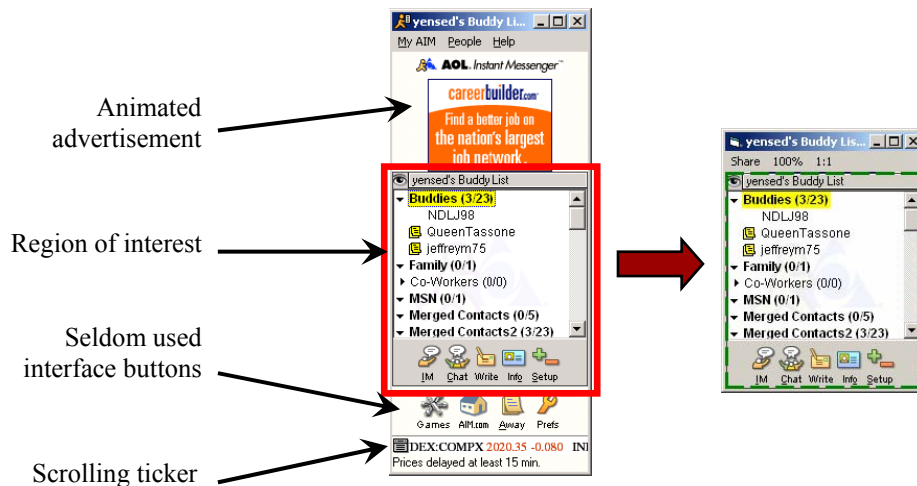


Figure 4.4. The most basic use of WinCuts is to specify regions of interest and get rid of all else. The user can regain scarce screen space as well as attention.

regions of interest and get rid of everything else. In Figure 4.4, the user makes a WinCut of the buddy list in an instant messaging client, eliminating the advertisements, seldom used interface elements, and a scrolling ticker. The WinCut of the instant messaging application not only occupies much less screen space, but is also much less distracting.

In fact, I quickly found that WinCuts were also useful for monitoring or peripheral awareness tasks. In these tasks, users are trying to keep abreast of updating information using the least amount of screen space possible. Using WinCuts, users can specify the ROI, scale it to an appropriate size, and move it to an appropriate location on the screen.

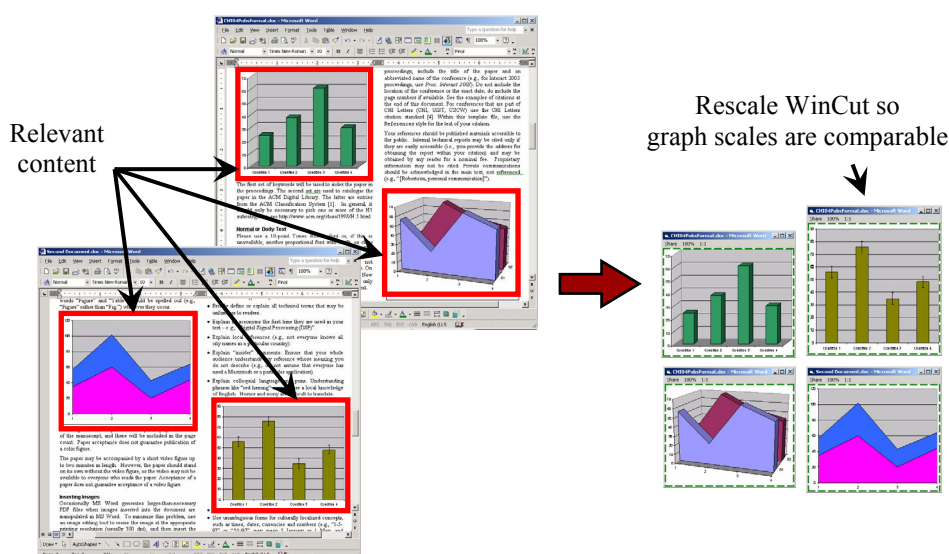


Figure 4.5. Users can make WinCuts of content from multiple sources, and then rearrange or even rescale them to make for easier visual comparison.

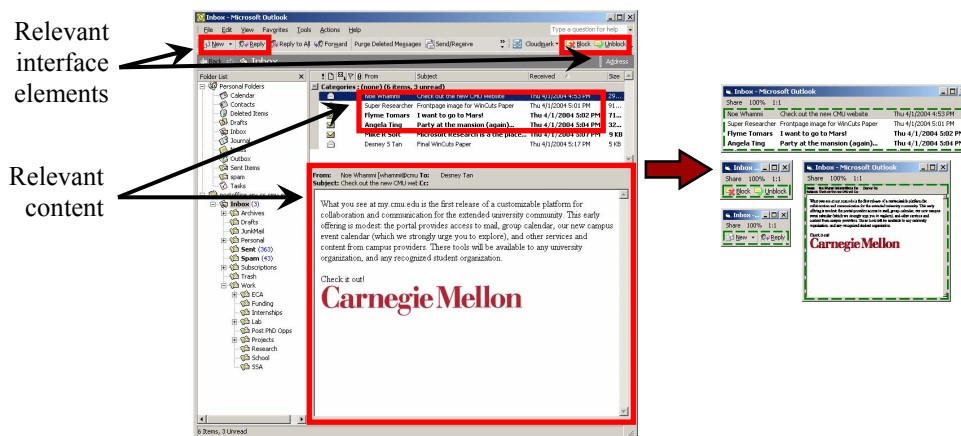


Figure 4.6. Using WinCuts, users can reconfigure entire interfaces or desktops.

A slightly different use of WinCuts is to organize useful content on the screen. In Figure 4.5, the user would like to compare the four graphs contained within two documents. If the user has a single display, they would have to flip back and forth between the two documents, bringing each into focus and viewing the graphs on that document before looking at the graph in the other document. If the user has multiple displays, they could lay the documents out side by side. While this might make it easier to compare the graphs, it is still difficult to lay the graphs out such that they are immediately next to each other. Information not relevant to this task is displayed in between the relevant information. With WinCuts, the user can make WinCuts of the relevant graphs and place them next to each other on the screen. Additionally, they can rescale the WinCuts such that axes on the graphs are equivalent, for even simpler comparison.

WinCuts are completely live regions and have no notion of whether they are capturing content regions or interface elements. Thus, in addition to organizing content layout, WinCuts are also useful for reconfiguring entire interfaces. Figure 4.6 illustrates one such scenario. Here, the user has spatially reorganized regions of the original e-mail client window to reduce the screen space used by this task. The user has reconstructed relevant interface components, including buttons on the toolbar. They have also chosen to display only three columns of the inbox, the sender, the subject, and the date, and only the last four messages that have arrived in the inbox. Finally, they have scaled the message pane down to occupy less space but to retain an overview of what is contained within messages. The user can interact with the inbox to change the message in the message pane as well as with the various interface components. They can also hit the 100% button on the message WinCut to read it at full size.

Extending this scenario, users may use WinCuts to completely reorganize their desktops in order to support a greater workflow. For example, a user authoring a specific section of a web page might make changes in an editor, publish the webpage, and then refresh and check the section in several different browsers. They might repeat this process many times to get the content correct. Doing this using standard management schemes would probably require that the user perform tedious windows management, bringing each to the front, issuing the refresh command, and then viewing the content. Instead, the user could make WinCuts of the refresh buttons and the relevant sections and lay them out accordingly so that the task becomes a matter of clicking on all the refresh buttons, which now exist next to each other, and viewing the changes.

Perhaps the most unexpected use of the system has been as a rapid interface prototyping tool. In fact, users can recreate entirely new interfaces by making WinCuts of various regions and scaling or rearranging them appropriately. Using this technique, I have explored how various interfaces would function if laid out differently. I have also explored several focus-in-context and fisheye view ideas by creating multiple adjacent or overlapping WinCuts and scaling them to different degrees. Without WinCuts, most of these ideas would have taken much longer to build and explore.

4.4.2 Multiple Machine Tasks

In small informal groups, users often come together with various pieces of information contained on their personal devices, such as their laptops. In current co-located collaboration scenarios, users have to view the contents of one machine at a time, for example when sharing a common projector. Alternatively, to view all the material together, users could either print out relevant material or combine it on a single machine. This is not optimal, especially when the material is dynamic and cannot be determined beforehand, requires live editing, or is interspersed within private information that the author does not wish to share.

Using WinCuts, users can easily exchange updating views of relevant information, which owners can continue to edit on their respective source machines. Furthermore, if there is a shared display available, such as a projector, multiple users can send their WinCuts to this shared visual space so that everyone has a consistent view of the shared information (see Figure 4.1). Using Visitor, users can manage WinCuts on the shared machine. Alternatively, if no shared screen is available, users can share information di-

rectly between their personal devices, sending WinCuts back and forth in a more ad hoc manner.

As currently implemented, WinCuts allows users to utilize display space offered by various devices on their desktops. For example, users can send WinCuts of peripheral tasks onto the screen of their laptop rather than taking up valuable screen space on their main machine. As I explore input redirection on remote machines, I expect that this will become an even more compelling scenario, as users can use WinCuts to take advantage of specialized input capabilities on various devices, such as pen input on the tablet PC. Also, this might be useful in placing particular parts of interfaces on a remote machine and interacting from afar, for example, using a laptop or PDA as a remote control for an application running on the desktop machine.

4.5 High-Level Implementation Details

I built the initial prototype of WinCuts as a standalone application in Windows XP using Microsoft Visual Basic and utilizing the Win32 Graphics Device Interface (GDI) API. This prototype has since been rebuilt with similar mechanisms using Microsoft C++ .NET. While the current implementation is specific to the Windows operating system, there is no reason why, given appropriate engineering effort, the general process could not be replicated across any operating system.

When a user specifies a desired WinCut, I first calculate the coordinates of the ROI within the source window. Next, I create a device context, into which I periodically force the entire source window to render, using the `printwindow` API call. From this device context, I perform a `stretchblt` to scale and copy the ROI into my WinCut. The reason I first do the `printwindow` is to ensure that occluded parts of the window that do not normally render are properly captured. In order to ensure that content remains relatively fresh, I periodically refresh the image for each WinCut. I am exploring schemes that dynamically update when necessary rather than being based on a timer. This would allow WinCuts to work at more interactive rates, but is technically difficult as it requires knowing when each individual application has repainted or needs to repaint any part of its window. To improve performance, subsequent WinCuts that depend upon the same source window reuse the appropriate device context that I have already created.

In order to redirect input, I currently activate the source window and bring it to the front when the cursor first moves into the WinCut content region. I then programmatically move the actual cursor to the corresponding spot on the source window. In order to simulate interaction with the WinCut, I draw a copy of the cursor on the corresponding piece of content in the WinCut. Hence, while all interaction actually happens directly on the source window, the user has all the feedback of interacting on the WinCut. One caveat to this approach is that source windows coming to the front can sometimes occlude information that is relevant to the task. I currently get around this by sending the source window to an extra display device that is not visible to the user. Another caveat is that since I do not explicitly handle them, popup menus and other windows that appear based on the location of the actual cursor appear on the source window and may not be seen on the WinCut. Solving these problems remains future work.

For remote WinCuts, I currently open peer-to-peer socket connections and send images of the printwindow device context, compressed as Portable Network Graphics (PNG), to the destination machine. I also send the corresponding coordinates so that appropriate WinCuts can be made on the destination machine. Sending subsequent WinCuts then, is as easy as sending additional coordinates along with the name of the device context.

4.6 Limitations of Current Implementation

The current implementation is meant to serve only as a proof of concept for the ideas behind WinCuts, and not as an end-product. In fact, the current implementation has several limitations. First, the output redirection mechanism is not optimal and slows down the system if WinCuts are made from too many source windows. The `printwindow` API call is an expensive operation, both in terms of processing as well as system memory. Since I periodically update the device contexts representing the contents of each source window whether or not the window has changed, this can start to be a drain on system resources. Currently users would notice a significant slow down on a 3.33 GHz machine that has 1 Gig of RAM when more than 12 average-sized source windows are trying to update. One solution to this would be to watch the windows event streams and to update only when the source window has changed. A better solution would be for the specific operating system to be built to natively support output redirection in a much more efficient manner.

A second problem is that since I currently operate at the GDI level in windows, certain streams of output such as video playback do not work in WinCuts. Additionally, even if they did, it would be difficult to send these at fast enough rates across the network and to other machines. I would like to explore mechanisms for dealing with media such as video and audio.

A third problem in the implementation is that input redirection is tricky in a system that implicitly assumes a single stream of input data. Currently, I perform input redirection in a somewhat roundabout manner, manipulating windows and the mouse pointer without the user knowing about this. Ideally, input redirection would be a cleaner mechanism, performed by replicating the input stream into a different window. Unfortunately, with the current system, doing this would lead to ‘focus thrashing’, with each input message causing different windows to fight for focus. Furthermore, since a window that is not in focus has no notion of active widgets, typing into a textbox that exists in an inactive window is surprisingly difficult. I do not see an immediate workaround with the current assumptions built into our operating systems. Perhaps a larger reconsideration of the way we architect these systems is in order.

4.7 Summary of Key Ideas

There are several key ideas that separate WinCuts from previous work and make it useful as a general mechanism for organizing and managing information across a variety of scenarios. First, users define WinCuts by specifying arbitrary window regions. It is important that users define these regions because it is hard for system designers to predict what a useful chunk of information is for every user and every task. It is also important that the region be flexible and not be tied to system representations, since these regions often do not mimic user perception. One exception of this is interface elements, which have fairly well-defined semantics. Realizing this, several systems and applications (e.g. Microsoft Visual Studio) already implement customizable and movable toolbars. Additionally, once created, WinCuts are tied to window regions, not screen regions. This allows users to continue to use their screen as they would like, moving, occluding, or minimizing the source window without affecting the content contained within WinCuts.

Second, WinCuts replicate regions rather than operate on the source window. The WinCuts system has little control over third party applications and windows and getting

them to do things they were not designed for such as rendering scaled content in multiple places and receiving multiple streams of input is an unnecessarily difficult task. Instead, WinCuts are copies of these regions that can be manipulated as independent windows in their own right.

Third, WinCuts use high-level I/O redirection to represent live content and permit users to interact with that content. One alternative implementation for a system such as WinCuts is to understand and access the low level semantics of each application on which it has to operate. However, this is tedious and not scalable as most operating systems do not place strict standards on how applications are implemented, leading to a myriad of factors to consider for each different application. Working at the I/O stream allows WinCuts, without having any knowledge or access to underlying applications, to operate across anything a user would normally be able to see and interact with on the screen.

Finally, WinCuts can easily be shared between machines. I have shown that this is useful for working across multiple machines, either in collaborative scenarios or in single user ones. While WinCuts is a useful tool when used on a single machine, it also allows us to reconsider our notion of what a computing unit is and to bridge the seams that currently exist between multiple physical machines and devices.

4.8 Future Work

I would like to examine the deeper implications of input redirection across machines. Technically, this could be solved by schemes such as turn-taking, floor control, or more complex conflict resolution. However, given that I am working within a fairly constrained social environment, I would like to understand and find the best set of interface and social solutions that would work for such environments.

Because of how easy and useful it is to create WinCuts, users usually end up with many more WinCuts than they had windows. Unfortunately, most windows management systems do not scale well to a large number of windows, some of which might be logically associated with others. With the growing number of windows presented on multiple display systems, researchers have already begun to explore methods that allow users to easily manage groups of windows as well as to perform simple operations on these groups (e.g. Smith et al., 2003). I will extend these systems to deal with even more ex-

treme numbers of windows, and hope that these systems will be useful even when users are not using WinCuts.

I would also like to explore the notion of device relevant WinCuts. For example, if a desktop system knows that the user is trying to create an interface on a remote device such as a PDA, perhaps it should provide the user with a proxy of the PDA screen on the desktop display so that the user can create the interface before sending it all at once to the PDA. This raises larger issues of how machines understand the specific affordances they provide as well as how they fit into a larger environment of other machines with different affordances.

Another area I am exploring is the utility of tying the ROI to the underlying information rather than window regions. In the current system, when a user scrolls the source window, multiple WinCuts may be affected since they are defined only by the geometric region of the window and not the semantic content. While this is useful in many scenarios, I realize that tethering WinCuts to actual content might provide additional utility. I would like to understand the tradeoffs associated with the added utility versus loss of generalizability across applications in doing this.

Finally, I plan to conduct formal evaluations to measure the usability of the interaction model surrounding WinCuts. I will also perform field studies deploying WinCuts to users involved in information work and group meetings as well as controlled studies to closely examine the usefulness of WinCuts in particular settings. I am still especially interested in how WinCuts might affect productivity as well as social interaction in co-located collaborative work.

“Public behavior is merely private character writ large.”

Stephen R. Covey

Chapter 5

Examining Social Issues when Working on Large Displays

In the previous chapters, I have examined the cognitive utility of physically large displays, showing how they can be used to improve performance on spatial tasks. I have also presented the WinCuts interaction technique, which was motivated by designing a tool to support co-located collaboration around physically large displays. In this chapter, I will explore some of the social issues, especially as they relate to privacy concerns, when users work on physically large displays.

Central to the creation of the Display Garden is the belief that there is an emerging trend in the workplace towards the use of large wall-sized displays, typically used in conjunction with traditional desktop displays. In some projects, such as Bishop & Welch’s (2000) Office of “Real Soon Now”, researchers have gone so far as to completely replace desktop displays with large-screen projection displays. Most of these researchers have observed that visitors treat information on these large displays as being public and do not hesitate to read or comment upon it.

A common explanation for this loss in privacy is the higher legibility of information presented on large displays. Because these large displays are typically viewed from a distance that is not proportionally scaled with the increase in display size, they often provide a larger visual angle, making them easier to see and read.

In this chapter, I present a novel paradigm for measuring whether or not a user has read certain content. Using this measure, I show that, even with constant visual angles and legibility, visitors are still more likely to glance over a user's shoulder to read information on a large wall-projected display than on a smaller traditional desktop monitor. I assert that, in addition to legibility, there are more subtle social factors that may contribute to the loss of privacy on physically large displays.

Following this, I explore one implication that the loss of privacy has on the way we interact with large touch screen displays. Many large touch screen displays employ alternative mechanisms for text input, including soft keyboards and handwriting recognition. These alternative input interfaces are intrinsically observable. That means that someone watching the typist use these interfaces can fairly easily reconstruct text that has been entered, an activity known as shoulder surfing. This is undesirable when typing any private text, but is especially problematic for passwords. I present the Spy-Resistant Keyboard, a novel approach to private text entry on large touch screen displays. I also present a study that shows the benefits of using such a keyboard.

5.1 Quantifying How Physical Size Impacts Privacy

One reason social phenomena revolving around privacy have not been well quantified is that there has not existed an easy way to measure if someone is looking at a particular piece of information, and to what degree they have cognitively processed the information. In this section, I describe the novel use of an implicit memory priming paradigm in order to quantify this phenomenon. I also show that people are more likely to peek at information, even private information, shown on a large display, even when legibility and visual angles are held constant.

5.1.1 Materials

I used two displays, an NEC MultiSync FE1250 22" monitor and a Sanyo PLC-XP30 LCD projector. Both displays ran at a resolution of 1024×768 and were calibrated to be of equivalent brightness and contrast. The image on the monitor was 16" wide by 12.5" tall. The image projected on a wall-mounted screen was adjusted to be exactly 66" wide by 49.5" tall. I set the displays up so that when either display was viewed from the participant's seat, the visual angle and the size of the retinal image would be identical (see Figure 5.1). I have already reported results in Chapter 3 suggesting that reading perform-




Figure 5.1. As with other experiments, I held the visual angle constant between the Small and Large Display conditions.

ance did not significantly differ on two such displays. Additionally, I ensured that someone using the system would not occlude any part of either displays from a participant.

I measured whether participants had read content on the displays with an implicit priming paradigm, usually employed to study learning without awareness (Schacter, 1987). In this paradigm, participants are presented with target words and are later tested, for example with stem completion, on their implicit memory for these words. In stem completion, participants are given the beginning of a word (e.g. mon__ for monkey) and must complete it with the first appropriate word that comes to mind. There are many possible completions for the letters “mon” (e.g. monarch, Monday, money, mongoose, monitor, monogram, monster, month, etc.) but priming is reflected by an enhanced tendency to complete stems with target words, words that had been previously seen.

In my study, I constructed seven e-mail subject lines and two e-mail messages that included a total of 30 target words selected from the Kucera and Francis (1967) norms. These words were selected to be between 5 and 12 letters in length and of medium frequency (mean: 87 per million). The initial three letters, or stem, of each word was unique to all other text given to the participant and each had at least 8 different completion possibilities. Additionally, I selected 33 filler words with the same criterion.

5.1.2 Procedure

Participants were seated in the armchair and handed a survey on multiple monitor preferences. They were informed that they would have exactly seven minutes to complete this survey and should look over their answers if they got done earlier. In their mind, this was their main, and only, task. The survey, consisting of one open answer and ten multiple-choice questions, was designed to take substantially less than seven minutes to complete.

I used a between-subjects design, with participants balanced by Gender and assigned randomly to one of the two Display Size conditions: small vs. large. While the participant completed the survey, the experimenter read the prepared target content on one of the two display setups. They viewed each e-mail message in the Microsoft Outlook e-mail client for three and a half minutes. The seven subject lines remained visible in the inbox for the entire seven minutes. The participant was video taped during the experiment.

After the survey, the participant performed the stem completion test. They were not informed of the purpose of the test. The test consisted of 63 stem completions: 3 practice questions, 30 filler questions, and 30 target questions. The 3 practice questions were followed by the 60 filler and target questions presented in random order. Following the test, participants completed a questionnaire explicitly asking whether or not they had read content on the experimenter's display while doing the survey.

5.1.3 Results

Twenty-four (12 female) intermediate to advanced windows users with normal or corrected-to-normal eyesight participated in this study. Participants ranged from 18 to 55 years of age (mean: 36.9). They received a software gratuity for participating.

I found significant differences between conditions in the number of stems completed with target words, suggesting that users had read more information displayed on the large display ($M=3.83$ words) than on the small one ($M=2.67$ words). This was true with both a loose metric that permitted different forms of the target words, as well as a concise one that allowed only exact forms that had been presented ($t(22)=2.0739$, $p=.04$; see Figure 5.2).

On post-test surveys, more users admitted to having read text on the display in the large screen condition (7 of 12) than in the small one (3 of 12), marginally significant by Fisher's Exact Test ($p=.089$). Additionally, video tapes showed users spending longer




Figure 5.2. Users responded with target words significantly more in the Large Display condition, suggesting that they had read text more in this condition.

periods, on average, viewing material on the large screen ($M=19$ seconds) than the small ($M=14$ seconds), though this difference was not significant.

5.1.4 Summary

In this section, I have described the novel application of an implicit memory test to measure whether a participant has read certain information on a given display. I have also presented results showing that, even with constant visual angles and legibility, visitors are still more likely to glance over a user's shoulder to read information on a large wall-projected display than on a smaller traditional desktop monitor. I assert that, in addition to legibility, there are more subtle social factors that may contribute to the loss of privacy on physically large displays. I believe that social convention prescribes that people have certain personal zones within which objects (information included) are deemed private. With few exceptions, any object outside of this zone is assumed to be public. Also, culturally, objects placed on walls are typically considered public. Most large wall-sized displays exhibit both sets of public cues.

In the following section, I will explore one implication that these findings have on the way we interact with large touch screen displays, namely that it is difficult to enter private text. I present the Spy-Resistant Keyboard, an interaction technique that provides a possible solution to private text entry even when typists are being observed. I also present a study showing the tradeoffs involved in using such a keyboard.



Figure 5.3. User typing with a soft keyboard on a publicly observable touch screen.

5.2 Secure Password Entry on Public Touch Screen Displays

Many multiple display environments are incorporating large touch sensitive displays as a key component. In fact, touch screens are becoming increasingly common, appearing on devices such as digital whiteboards, tablet PCs, as well as ATM and debit card machines. Many of these devices assume that the touch screen is the primary input mechanism and make using traditional text input mechanisms such as physical keyboards inconvenient. As a result, these devices often employ alternative mechanisms for text input, including soft keyboards and handwriting recognition.

The soft keyboard functions like a hardware keyboard except that users touch an on-screen image-map to type (see Figure 5.3). With handwriting recognition, users enter text by writing on the touch screen. These alternative input interfaces are intrinsically observable. That means that someone watching the typist use these interfaces can fairly easily reconstruct text that has been entered, an activity known as shoulder surfing. This is undesirable when typing any private text, but is especially problematic for passwords, especially given my results implying the intrinsic loss of privacy on large displays.

Users typing on these public touch screen displays can take precautions that make it harder for a casual observer to obtain their password. For example, they may physically obscure the display so that observers cannot see the interface feedback or the results of their actions. However, this can be difficult on touch screens that are placed in locations that make blocking it inconvenient or that are larger than the user can physically block with their body. Users may also adopt other strategies, such as quickly adding and deleting characters that are not in the intended password in order to confuse observers. Unfor-

tunately, my observations show that this strategy usually leads to increased levels of mistyped passwords and does not add much security against alert observers. Additionally, taking intentional actions to protect one's password may be socially awkward as it conveys a lack of trust in the observers.

In my work, I set out to design an onscreen virtual keyboard that ensures security from shoulder surfers without requiring typists to take explicit precautions. In fact, I assume that observers can openly watch any part of the typist's interaction with the keyboard. Hence, the keyboard must not only ensure that the typist's actions cannot be easily converted into knowledge of the password, it must also prevent observers who do not explicitly know the password from repeating the typist's actions in order to enter the password.

In this section, I present a novel approach to designing keyboards for entering private text on large public touch screen displays. This approach introduces indirection by utilizing an auxiliary mapping that allows typists to focus their attention on a particular part of the keyboard, while observers have to pay attention to and memorize the entire keyboard. I describe one particular instantiation, which I call the Spy-Resistant Keyboard, as well as the main design decisions leading to its development. I present results from a user study evaluating both the usability as well as the additional security offered by the Spy-Resistant Keyboard. Finally, I discuss future work extending these ideas.

5.2.1 Background

In many computer systems, users have to authenticate themselves to access sensitive data and services. There are two basic components to any user identification scheme: the interaction with which the user implicitly or explicitly provides identifying data to the system; and the ensuing check by the system to ensure that this data matches some prior piece of information it has about the user. There exists a large body of work on the latter, examining possible attacks (Neumann, 1994) as well as how to protect against them as the system either sends the data to be authenticated, or does the comparison locally. For a review, see Halevi and Krawczyk (1999).

Technical issues, while important, are not the only component of secure authentication. In fact, Hitchings (1995) asserts that treating security as a purely technical issue has led us to produce mechanisms that are less effective than they could and should be. Davis

and Price (1987) add that since security necessarily involves people, human factors should be carefully considered in designing effective security mechanisms. In my work, I focus on creating interaction techniques that make it more secure for users to provide private data to the system.

Currently, users have three basic methods to provide private data and authenticate themselves: tokens, biometrics, and private knowledge such as passwords. Token-based methods utilize something a user possesses, such as an identification card, to verify their identity (Brostoff & Sasse, 2000). Such methods often require costly construction and distribution of tokens, as well as installation of specialized sensing hardware. Additionally, possession of a token does not necessarily imply ownership, and theft or forgery is a serious threat to these systems.

Biometric methods identify individuals based on distinguishing physiological or behavioral characteristics. These methods include signature, keystroke pattern recognition, voice, vein geometry, as well as eye-based, facial, finger, and palm imaging. For a detailed review, see (Jain, Hong, & Pankati, 2000). Just as with token-based methods, biometric methods involve costly hardware, and characteristics can be stolen or forged. Furthermore, since these characteristics cannot be easily replaced, theft is more costly than it is with other methods.

The third class of methods, which remains dominant on many computing systems, verifies access privileges with pieces of knowledge such as passwords known only to the user. Historically, the choice of passwords has been such a prevalent problem that the National Institute of Standards and Technology (1995) has published a document advising users of proper password selection and use. They recommend picking random strings of characters and keeping different passwords for different accounts. Unfortunately, this places a large cognitive strain on users, who have to remember an increasing number of passwords (Adams, Sasse, & Lunt, 1997).

To alleviate this problem, various researchers have proposed alternatives and augmentations to standard text passwords. For example, numerous researchers have proposed graphical passwords, in which users have to select predetermined positions within an image (Blonder, 1996), recognize images such as faces (Brostoff & Sasse, 2000), or sketch drawings that are recognized by the system (Jermyn, Mayer, Monroe, Reiter, & Rubin,

1999). Others have explored the use of cognitive passwords, which use a question and answer session either prompting for personal details or word associations (Zviran & Haga, 1990) known only to the user.

All these knowledge-based methods make one fundamental assumption, that no one other than the user knows the password. In fact, these methods can only be as secure as the user's ability to keep the password secret. With the introduction of large touch screen displays that utilize onscreen soft keyboards, learning someone's password has become as easy as shoulder surfing, or watching them type it in. In fact, even apart from adversarial observers, I have shown in the previous section that people are generally more likely to peek at private content on large public displays, making unintentional viewing of password entry more likely than before. The use of one-time passwords (Haller, Metz, Nesser, & Straw, 1998) is the closest method I have found that might protect against such attacks. However, this method usually requires that users constantly learn new passwords, which is impractical for most computer users. Also, it cannot be applied to generic private text entry.

In my work, I aim to design a general purpose interaction technique and interface that allows users to enter private text such as passwords on public touch screen displays without the risk of revealing them to observers.

5.2.2 Design Approach

My approach to designing virtual keyboards that add security against shoulder surfing involves breaking the typing interaction into two phases, the mapping phase and the selection phase.

In the mapping phase, the keyboard presents the typist with some method of uniquely mapping each character to some property, which serves as an alternate representation of the character. This property could be anything ranging from another character to a color to the shape of the button to the spatial location of the button, and so on. While each keyboard should use as simple a mapping as possible, it should always re-randomize the mapping after each character is typed so that observers cannot learn the mapping over time. The typist locates the character they wish to type and mentally notes the specific property that represents this character. They need not focus any attention on properties

representing other characters. Once they have done this, they signal to the keyboard that they are ready to move into the selection phase.

In the selection phase, the keyboard removes the mappings, usually by blanking characters from the keyboard. The typist completes their interaction by specifying the property that represents the character they wish to type. They then repeat this process for each character in their password.

5.2.3 Justification of Approach

This approach provides two mechanisms that make it hard for the observer to derive the character that has been typed by watching the actions of the typist. First, since the keyboard is randomized for each character, there is no way for the observer to repeat the typist's actions to reproduce the correct password. Second, the explicit two-phase interaction conceals the typist's intention until pertinent mapping information is hidden. This makes it difficult for the observer to derive the character typed even when they can watch the entire interaction. In fact, the observer has little information to help focus their attention in the mapping phase, and is forced either to guess which characters they should focus on or to memorize mappings for all keys. The former is fairly unreliable in figuring out what has been typed and the latter is difficult without using recording equipment.

5.2.4 Spy-Resistant Keyboard

I instantiated this approach in an interface I call the Spy-Resistant Keyboard. This keyboard randomizes the spatial location of all characters as each password character is entered.

The Spy-Resistant Keyboard is composed of 42 Character Tiles, two Interactor Tiles (labeled "Drag Me..."), a textbox for feedback, a backspace button, and an enter button (see Figure 5.4). Each Character Tile is randomly assigned a lowercase letter, an uppercase letter, and either a number or a symbol, all positioned vertically on top of each other. Lowercase letters are always on the top row of each tile and have a red background; uppercase letters are placed in the middle and have a green background; numbers and symbols are positioned on the bottom and have a blue background. Since there are exactly 42 numbers and symbols combined, but only 26 letters, some letters are repeated. Just as each button on a standard keyboard represents two characters, depending on the state of the caps lock or shift keys, each Character Tile represents three characters, depending on



Figure 5.4. The Spy-Resistant Keyboard. . In the first phase of typing, the mapping phase, the user must first find the character they would like to type and note its location. For example, if they are trying to type the capital letter “Z”, they would scan the three green rows, finding the letter on the seventh tile on the second row.

the state of shifting. Rather than having a fixed shift state for the entire keyboard, as traditionally done, each tile has a randomly assigned shift state, indicated by the red line under the active character.

In order to select a character on the Spy-Resistant Keyboard, the typist first locates the tile that contains the character to be typed. They remember the mapping by noting the location of this tile. Next, the typist clicks on one of the Interactors at the bottom of the keyboard to cycle through shift states and move the red underline to the desired character (see Figure 5.5). Clicking on the Interactor moves the underline to the next character on each tile. Note that since the underlines start on different types of characters on each tile, the underline will eventually cycle through all characters on each tile.




Figure 5.5. Each Tile on the keyboard begins with a random shift state, indicated by the red underline. Hitting the Interactor moves the shift state on all Tiles. In our example, the user taps the Interactor to cycle through states and get the underlining under the letter “Z”




Figure 5.6. To complete the selection, the user drags the Interactor to the appropriate Tile. It is very difficult for an observer to reconstruct what has been typed (a letter “Z” in this case).

knowing that the typist has clicked on the Interactor but not knowing which tile the typist is focused on gives the observer no useful information about the kind of character being typed.

Finally, the typist drags the Interactor towards the Character Tile on which the desired character resides. Upon the start of the drag interaction, the system knows that the user has located the character and moves into the selection phase, blanking all Character Tiles (see Figure 5.6). Hence, without knowing where the Typist is going to drop the Interactor, adversarial observers have to memorize the location of all characters on the keyboard so that they can reconstruct the typed character from the location of the drop. I anticipated that this would be a very difficult task, if not impossible. Each tile highlights as the typist drags over it. The typist drops the Interactor on the desired tile and the character is entered. The keyboard re-randomizes characters and the typist repeats the process to select the next character. After beginning the drag, the typist may also drop the Interactor on anything other than a Character Tile to reset the board and get a new set of characters, in case they lose track of their target.

5.2.5 Design Rationale

The Spy-Resistant Keyboard is a product of iterative testing and design. Specifically, it has evolved in three regards: the basic interaction mechanism, the design and layout of tiles, and the mechanism for specifying the shift state.

5.2.5.1 *Basic Interaction Mechanism*

I had initially used a tapping gesture for the interaction mechanism. In early prototypes, typists first tapped the Interactor to indicate that they had located the desired character and were ready for the tiles to be blanked. They then tapped on the appropriate tile to type that character. Most typists using this prototype quickly realized that they were much faster using one hand to tap the Interactor and the other to tap the tile. Unfortunately, most typists would anticipate their action by positioning their typing hand directly over the tile to be typed even before the tiles were blanked. This gave observers the opportunity to look at the character before it was blanked, negating any benefit of using the interface. Rather than having typists consciously sequence their actions, I decided to build this requirement into the interaction technique. The current drag-and-drop mechanism provides an intuitive interaction that forces the user to perform the two phases one at a time, first blanking the tiles by starting to drag the Interactor and then selecting the desired character by dropping it.

5.2.5.2 *Design and Layout of Tiles and Interactors*

The layout of Character Tiles significantly affects how quickly users are able visually search the keyboard and find specific characters. In early prototypes, tiles were designed with the set of three characters running across each tile. This meant that similar characters (i.e. lowercase, uppercase, as well as numerals and symbols) were grouped in vertical rows. Performing the visual search with this design required typists to scan across all tiles, stopping to look at every third character. Alternatively, the typist could scan the characters from top to bottom, looking only at every third column. Unfortunately, early observations showed that this was an extremely difficult task, and even the most practiced users had trouble finding characters.

Since we are more accustomed to scanning contiguous blocks of text running from left to right, grouping letters horizontally by stacking the three characters on each tile made the search task easier. Also, color coding each type of character aided in identifying the appropriate rows to scan. These adjustments are consistent with suggestions laid out by Wickens and Hollands (2000) in their design recommendations for directing attention in display space. Furthermore, I explored several different configurations of these tiles, for example a single row of 42 columns or 2 rows of 21 columns. I settled upon the current layout of 3 rows of 14 columns, as it provides a relatively small number of rows but maintains an aspect ratio that fits nicely on most displays.

Another layout decision I made was to include two Interactors, one on either side of the keyboard. Although this was initially done to make the interaction comfortable for both right and left-handed users, I found that both groups seem to use both Interactors interchangeably, depending on where they are standing in relation to the interface and where the target character lies. I assert that this is a useful design element to include in large, touch screen displays such as the SMART Board™.

5.2.5.3 Specifying Shift State

Finally, I had initially used one tile to hold each character, eliminating the need for a shift key. However, since I wanted to separate the different character types to help users constrain their visual search, knowing the layout of the keyboard as well as where a typist dropped the Interactor would necessarily reveal the kind of character entered. While this information does not reveal the entire password, it drastically reduces the search space of all possible passwords. I could have also used a single shift state for the entire keyboard, much as traditional soft keyboards do. However, since this state would have to be visible, knowing that the typist had hit the Interactor to change the shift state would also reveal the kind of character typed. Hence, I decided to randomize the shift state for each character. This adds little additional load on the typist as they still only have to focus on one particular tile and its shift state, but adds significant complexity for the observer who has to monitor all possible shift states in addition to all possible characters.

5.2.6 User Study

I compared the Spy-Resistant Keyboard to a standard soft keyboard in order to examine usability as well as additional security it provides. To ensure equivalent visibility by observers, I used the same font, 16-point Courier bold type, for all characters in each interface.

5.2.6.1 Participants and Setup

Six pairs of Microsoft employees (8 males, 4 females) volunteered to participate in the study. All users had normal or corrected-to-normal eyesight, and all were right-handed. The average age of users was 28.8, ranging from 21 to 38 years old. Users received a small gratuity for participating.

I ran the study on a SMART Board™ 3000i, which provides a physically large rear-projected touch screen display. The display was approximately 53" tall by 40" wide and

ran at a resolution of 1024×768 . Users stood in front of the display and interacted with the interfaces by touching the display with their fingers (see Figure 5.3).

5.2.6.2 Task and Procedure

Before beginning the test, I gave users paper-based instructions on how to type with the soft keyboard as well as with the Spy-Resistant Keyboard. Both users took turns practicing each interface by typing in a password I provided. All users were able to complete each practice password in less than two and a half minutes.

For each trial in the test, one user played the role of Typist while the other was the Observer. The Typist used one of the two interfaces to type in passwords. The Observer watched the Typist to discover the passwords. Typists were allowed to use any technique they wished to prevent the Observer from figuring out the password. However, in order to simulate public visibility of the display, they were not allowed to explicitly physically obstruct the Observer's view of the keyboard.

Observers were also allowed to use any technique they wished to watch the Typist and figure out the password. For example, they could move around to get the best view of the screen and many took notes to help them reconstruct the passwords. After each entry, the Observer recorded what they thought the password was. The pair performed each entry twice for each password.

5.2.6.3 Design

I assigned each Typist one easy, one moderate, and one difficult password for each interface. All passwords were 8 characters long. I randomly chose the easy passwords from the set of English words having Kucera-Francis (1967) familiarity and concreteness ratings between 300 and 700 (e.g. *contract*). I chose the moderate passwords to contain 3 to 5 letter English words surrounded by random characters (e.g. *#back\$Jr*). The difficult passwords were completely random sequences of 8 characters (e.g. *s%g7^Lp=*).

I used a 2 (Interface: Soft Keyboard vs. Spy-Resistant Keyboard) \times 3 (Password: Easy vs. Moderate vs. Difficult) within-subjects, dyadic design with repeated measures. Each user performed each of the 6 conditions twice, once as the Typist and once as the Observer. I balanced the order of Interface across pairs, with each member of a pair using the interfaces in the same order. I randomized the order of Password.

I collected the following dependent measures from the Typist in order to compare usability of the two interfaces: completion time, number of backspaces, and error rates for each password entry. In order to determine the level of security provided by the interfaces against watchful observers, I collected the Observer's guesses from each password entry. Finally, users filled out a post-test questionnaire indicating their preference for each of the interfaces.

5.2.6.4 Results

5.2.6.4.1 Typist Performance: Usability

I analyzed the average completion time required to enter each password with a 2 (Interface: Virtual Keyboard vs. Spy-Resistant Keyboard) \times 3 (Password: Easy vs. Moderate vs. Difficult) repeated measures analysis of variance (RM-ANOVA). I found a significant main effect of Interface ($F(1,11)=114.11$, $p<.0001$), with the Soft Keyboard resulting in faster completion times on average (see Figure 5.7).

I found no significant difference in the number of backspaces hit for each of the conditions. In fact, Typists seemed to hardly ever use the backspace key (average of about 1 backspace hit every 20 passwords typed). I also found no significant differences in the error rate for entering passwords. In fact, only 9 out of a total of 144 passwords were entered incorrectly, and most were off by a single character.

5.2.6.4.2 Observer Performance: Security

I compared each guess made by the Observer to the password typed and generated two metrics representing the level of security offered by the interface: a strict metric, the




Figure 5.7. Main effect of Interface for average time to type each password.




Figure 5.8. Main effects of Interface for the level of security, measured by errors in guessing the password

number of characters in each guess that did not match its typed counterpart exactly; and a loose metric, the Levenshtein (1966) distance, or number of deletions, insertions, and substitutions required to transform the guess into the typed password. This loose metric accounted for characters that were shifted in position. Both these metrics produced ratings on a scale of 0 to 8, with 0 indicating poor level of security and 8 indicating strong level of security offered by the interface.

I performed similar 2×3 RM-ANOVAs for the level of security offered by the interfaces. This analysis revealed a significant main effect of Interface for both the strict metric ($F(1,11)=641.47$, $p<.0001$) as well as the loose one ($F(1,11)=1250.68$, $p<.0001$), with the Spy-Resistant Keyboard resulting in far stronger security, on average. In fact, although most observers were able to fairly accurately guess entire passwords with little error on the Soft Keyboard, they were not able to get even one of the eight characters correctly with the Spy-Resistant Keyboard. Additionally, the loose metric revealed a significant main effect of Password ($F(1,11)=7.31$, $p=.004$), with progressively higher security using the more difficult passwords (3.85 vs. 4.31 vs. 4.44, on average, for Easy, Moderate, and Difficult passwords respectively). These results, illustrated in Figure 5.8, indicate the drastically improved level of security offered by the Spy-Resistant Keyboard against shoulder surfers.

5.2.6.4.3 Subjective Ratings

In addition to performance data, I gathered user preference data on 5-point Likert scales after the study. Users found the Soft Keyboard ($M=4.92$) significantly easier to use than the Spy-Resistant Keyboard ($M=2.42$), ($t(11)=16.58$, $p<.0001$). However, users indicated that they were also significantly less comfortable with using the Soft Keyboard to enter their passwords ($t(11)=-13.01$, $p<.0001$, $M=1.17$ vs. $M=4.50$, Soft Keyboard vs. Spy-Resistant keyboard respectively). This sentiment was further supported by users feeling like they had much more difficulty acquiring useful information when observing someone using the Spy-Resistant Keyboard ($M=4.50$) as opposed to the Soft Keyboard ($M=1.67$), ($t(11)=-10.47$, $p<.0001$). Additionally, most users agreed that the extra security was worth the extra effort, especially since most passwords are relatively short. This is important since it has been shown that security measures that are not compatible with user perceptions often end up being circumvented, thereby undermining system security (Adams, Sasse, & Lunt, 1997).

5.2.7 Discussion and Future Work

Study results suggest that the Spy-Resistant Keyboard imposes a tradeoff between efficiency of entering text and the security of text entered. Using the Spy-Resistant Keyboard takes about twice as long as a soft keyboard, but distinctly makes the perceived as well as actual level of security provided against observers much stronger.

In future work, I will explore schemes to make the selection task easier, while still maintaining similar levels of security against observers. One improvement to the current keyboard would be to make remembering the location of a tile easier by providing landmarks within the keyboard. These landmarks could simply be spaces in between sets of tiles within each row, or they could be more complex background images or tiles of different shapes. While this would make selection of the character easier, this would not significantly speed up the visual search task.

One promising alternative is to completely eliminate the visual search task by not initially randomizing the characters on the keyboard. Instead, I would lay the keyboard out as it is normally displayed. However, when the user starts the drag, I would hide the characters and then animate each key into a new position, thus providing a one-to-one mapping of a character to a new spatial location (see Figure 5.9). With this mapping, the typist would have to watch the changing location of one key, but the observer would have




Figure 5.9. (left) We can eliminate the visual search task by laying out the keys as they normally would be. (right) The mapping is performed by blanking the keys and animating them into new positions. After mapping, the sequence of keys from left to right, top to bottom, is 9,4,8,*,2,0,5,7,1,#,3,6. To type a character, the user specifies its new position.

to know where all keys started and ended in order to later reconstruct what has been typed. Additionally, this interface does not need an explicit shift state because of the way randomization happens after the user has found the character. In fact, this allows all characters to be displayed at once and grouped in whatever manner is most convenient.

I found in the study that observers who devised strategies either tried to monitor the typist's gaze or concentrated on only a small region of the display hoping that the desired character lay there. Although these may be more effective than other strategies, it would still take many observations before gaining access to the full password. In future work, I will explore schemes that provide feedback on the remaining safe lifetime of a password based on the number of times it has been entered and the types of interfaces used. In such a scheme, a password's level of safety would decay much more rapidly when entered on a public touch screen keyboard than on a private desktop machine in the user's personal office, and users would be warned accordingly.

Finally, I must stress that the Spy-Resistant Keyboard does not do well to protect against observation that may be rewound and replayed, for example from an observer recording with a video camera. Techniques I have devised so far that protect against this kind of attack have relied upon compound passwords, which are combinations of two or more properties. In future work, I plan to further explore these techniques.

5.2.8 Summary

In this section, I have presented a novel approach for designing virtual keyboards that protect typists from revealing private text to watchful observers. By breaking the typing interaction into two distinct phases, I provide a level of indirection that allows the typist to focus on a specific part of the keyboard while the observer has to memorize the entire keyboard to reconstruct the character typed. I described the Spy-Resistant Keyboard, one instantiation of such a keyboard, as well as important design decisions made in building this interface. I ran a study showing that although the Spy-Resistant Keyboard was slower to use, it provided a level of security significantly higher than that of a traditional soft keyboard. The study also showed that users thought the extra level of security was worth the additional effort. Finally, I presented several ideas that have the potential to speed up typing on such a keyboard while maintaining similarly high levels of security.

"Engineering is not merely knowing and being knowledgeable, like a walking encyclopedia; engineering is not merely analysis; engineering is not merely the possession of the capacity to get elegant solutions to non-existent engineering problems; engineering is practicing the art of the organized forcing of technological change... Engineers operate at the interface between science and society..."

Dean Gordon Brown

Chapter 6

Exploring Pragmatic Issues in Large Display Systems

6.1 Introduction

Thus far, I have described work done to understand and exploit affordances of physically large displays. However, we must also remember that the task of integrating all our understanding and tools to create rich computing environments with multiple display systems is intrinsically an engineering effort. In setting up the Display Garden, I encountered and solved interesting technology-related problems that come with working on physically large displays. In this chapter, I discuss two areas revolving around pragmatic issues of integrating large displays into our workspaces.

First, since I was front projecting to create my large displays, users in the room often found themselves working between the projector and the display surface. This caused undesirable projection on the user as well as temporary blindness from looking into the bright light of the projector. To alleviate this problem, I developed Pre-emptive Shadows, a technique that uses a camera-projector system to detect and turn off pixels that would otherwise be needlessly cast upon users' bodies and faces.

Second, having multiple displays in the Display Garden enlarged the physical display area, allowing the system to present information across much wider visual angles from the user. Also, since displays were placed at different depths or framed by physical bez-

els, physical discontinuities were introduced in the presentation of information in the workspace. Yet, relatively little is known about the how to best present information to the user given these display characteristics. I describe an experiment that utilizes a divided attention paradigm to explore the effects of visual separation and physical discontinuities when distributing information across multiple displays. Results show reliable, though relatively small, detrimental effects when information is separated within the visual field, but only when coupled with an offset in depth. Surprisingly, physical discontinuities such as monitor bezels and even separation in depth alone do not seem to affect performance on the set of tasks tested.

6.2 Pre-emptive Shadows

As discussed in the related work chapter, spatially immersive displays that physically surround the viewer with a panorama of imagery are becoming common. Many of these displays are room sized, accommodate several viewers, and implemented with several fixed projection display units. Since rear projection requires a large amount of space, many systems utilize front-projected displays. Users interacting with these displays often occlude the light from reaching the display surface. This has the dual effect of (1) casting shadows on the display surface, and (2) projecting undesirable, and often blinding, light on the user (see Figure 6.1).

Researchers have done work to eliminate shadows by using multiple redundant projectors placed at extreme angles to ‘fill in the blanks’ with pre-warped images (Sukthankar, Cham, & Sukthankar, 2001; Summet, Abowd, Corso, & Rehag, 2003). These

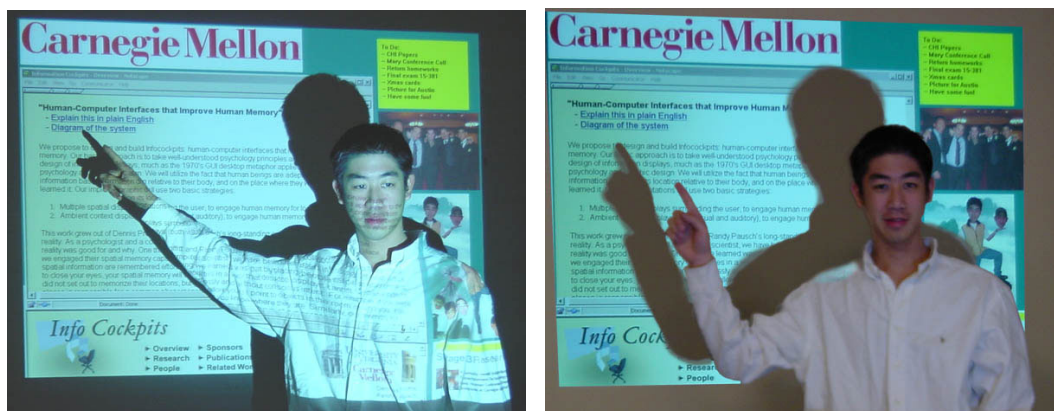


Figure 6.1. (left) Blinding light shining on user standing between projector and wall. (right) Blinding pixels turned off using Pre-Emptive Shadows.

techniques require fairly large hardware setups consisting of cameras and multiple projectors distributed over a fairly large space. They are also only starting to work at interactive rates. In many applications, the shadows go largely unnoticed, while the projection on the user causes unwanted distraction. For example, while presenting to an audience, a speaker might move in front of the projector, both causing undesirable projection on the body, as well as temporary blindness from the bright light.

In my system, I attach a camera to the projector and locate the shape of the user relative to the display surface. I use this information to turn off the pixels that would have been needlessly cast upon users, blinding them and creating shadows on the display surface (see Figure 6.1). I demonstrate the effectiveness of this technique by measuring the amount of light that reaches the user's eyes with and without this system.

6.2.1 Implementation

The initial implementation of my camera-projector system used background differencing to detect the shape of the user. Since I know the image that is projected at each instant, such a system merely has to check for pixels that deviate from the anticipated values. This approach had several problems. First, the physical offset of the camera and projector made matching pixels seen by the camera to the corresponding projected pixels a non-trivial task. This problem is aggravated by the drastic lighting changes caused by dynamic content (e.g. movies). Second, choosing the appropriate thresholds to separate pixels that were hitting the intended surface from those that were not was difficult when parts of users and their clothing mimicked the reflective properties of the display surface.




Figure 6.2. User casts IR shadow on camera lens. This matches the projected shadow and is used to turn off the appropriate pixels.

Thus, rather than illuminating users from the front and trying to identify them by reflective properties, I uniformly light the display surface with non-visible IR light from the rear and identify users using occlusion. In the current system, I attach an infrared (IR) camera to the projector. I place IR lights as close to the display surface as possible (~1 foot), lighting the surface while reducing the possibility of illuminating the user. Users, backlit by the non-intrusive IR light, cast a robust shadow onto the lens of the camera (see Figure 6.2 for setup). I assume that our camera and projector share a common focal point. Since I am not doing any per-pixel differencing, it is sufficient that the IR shadow cast on the camera closely matches the shadow that would be cast by the projector on the wall. I process this image using standard machine vision techniques provided by the Intel Image Processing library and use this information to mask out the appropriate projected pixels. Running on a PIII 700 MHz computer, the system tracks multiple objects at 30 frames per second.

6.2.2 Evaluation

To quantify the effectiveness of Pre-emptive Shadows in reducing blinding light, I compared the difference in brightness with and without the system. The measurable quantity that most closely corresponds to brightness is luminance, or the intensity of light per unit area of its source (Ryer, 1997).

Luminance (cd/m ²)	Condition
1	Dark room
130	White Paper in good reading light
4,000 – 6,500	<i>With pre-emptive shadows</i>
8,200	Florescent Lamp
65,000 – 120,000	<i>Without pre-emptive shadows</i>
150,000	Threshold of visual tolerance

Figure 6.3: Luminance values for various conditions.

I used a Sekonic L-508 Zoom Master photometer to make luminance measurements. I measured the luminance values for a user standing 8 feet away and looking directly into an Epson Powerlite 703c projector, rated at 1000 lumens. I found that the luminance with Pre-emptive Shadows was about an order of magnitude less than without (see Figure 6.3). This roughly corresponds to a 5-fold perceived difference in brightness. More impor-

tantly, the light level is reduced from being close to the threshold of tolerance to more comfortable levels.

6.2.3 Summary

Pre-emptive Shadows is a method that eliminates unnecessary projection that is cast upon the user rather than the display surface. I have shown that the perceived light hitting the user in front of the projector is reduced from excessive, and painful, levels to tolerable levels. One interesting idea might be to apply Pre-emptive Shadows to the domain of creative stage lighting. Inverting the current approach, we could develop a system that tracks and illuminates only the actors without throwing the spotlight on background elements.

6.3 Effects of Separation and Discontinuities Across Multiple Displays

Aside from spatially immersive displays, multiple integrated displays have long been used in environments that require multiple people to simultaneously monitor and interact with complex visual information. As discussed in the related work chapter, such environments include control rooms, operations centers, trading floors, and planning rooms. Recently, there has been a trend in the marketplace towards similar multiple display systems in more traditional workspaces. In both cases, having multiple displays enlarges the physical display area, allowing the system to present information across much wider visual angles from the user. Also, since displays are often placed at different depths or are framed by physical bezels, physical discontinuities are introduced in the presentation of information in these workspaces. Yet, relatively little is known about how to best present information to the user given these display characteristics. In this section, I describe a study designed to explore the effects of visual separation and physical discontinuities when distributing information across multiple displays.

To isolate and understand individual factors of interest, I created a display system that allowed me to carefully control the separation and discontinuities associated with multiple displays. I ran a study that utilized a divided attention paradigm across several different display conditions. The test included a primary task done in conjunction with a secondary or tertiary task. In the primary task, users had to proofread and identify grammatical errors within a set of text articles. While doing this, users also performed the secondary task, notification detection. In this task, users had to detect and act upon visual

changes outside the focal region of the primary task. Upon detecting notifications, users performed the tertiary task, text comparison, in which they had to cross reference and compare content displayed in multiple locations on the displays. I picked these tasks to be representative of tasks information workers perform while multitasking in a single user desktop situation.

Results from the study demonstrated a reliable, but small, detrimental effect on performance from separating information within the visual field, but only when it is further separated by depth. I found that physical discontinuities introduced by bezels or depth alone had no effect of performance for my set of tasks. I conclude with design recommendations.

6.3.1 Background

Swaminathan and Sato (1997) summarize much of the multiple display system work by describing three distinct multiple display configurations: (i) distant-contiguous configurations consist of multiple displays placed at a fairly large distance from the user so as to occupy the same visual angle as a standard desktop monitor; (ii) desktop-contiguous configurations consist of multiple displays placed at a distance equivalent to a standard desktop monitor so as to drastically widen the available visual angle; (iii) non-contiguous configurations consist of display surfaces at different distances from a user and that do not occupy a contiguous physical display space.

Most systems I have examined, and certainly the Display Garden, fall into the latter two categories. These systems share one characteristic: information is displayed across a wider visual field such that not everything is always contained in the foveal region. In fact, this is true even of many traditional desktop and distant-contiguous systems, since the typical visual angle of a display is 20-40 degrees, while foveal vision covers only about 2 degrees. In addition, non-contiguous configurations introduce physical discontinuities as information is separated at different depths or by physical objects. In my work, I explore the effects of visual separation and physical discontinuities when distributing information across multiple displays.

6.3.1.1 Human Vision and Peripheral Information

There has been a long history of work in psychology and psychophysics documenting the size and shape of the visual field. In their work, Carrasco & Naegele (1995) present the

eccentricity effect, which shows that targets presented near the point of visual fixation are noticed much more easily than targets further away. Wolfe, O'Neill, and Bennett (1998) present a summary of visual explanations of this effect as well as a new explanation claiming that attention is partially modulated by eccentricity, leading to higher activation and faster search times for nearer objects. Additionally, they show that these eccentricity effects are reduced when there are fewer distractions on the screen. Other researchers have shown that mental workloads greatly affect the size and shape of the visual field. For example, Rantanen and Goldberg (1999) show that heavier workloads not only shrink the visual field by up to 14%, but also cause it to be vertically shorter and horizontally elongated.

Researchers, aware of the capabilities of the human visual system, have designed various tools that leverage peripheral vision and attention. For example, Cadiz, Venolia, Jancke, and Gupta (2001) provide a wide range of awareness information on the side of the display in their Sidebar system. In their work, they build upon previous research investigating methods of providing the most peripheral information while having the least impact on main task performance (Maglio & Campbell, 2000; McCrickard, Catrambone, & Stasko, 2001). Grudin (2001), in observing how users use multiple displays, asserted that the division of space afforded by multiple non-contiguous displays is sometimes beneficial over having a single contiguous space. He explains that the divisions created often help users segment the working space not only to “park objects out in the periphery” but also to more effectively assign specific functions to each subspace. Given this assertion that the physical divisions seem to create separate mental subspaces, I expected the divisions to be a distraction and to add cognitive load when a task was split across two of these subspaces.

In addition to eccentricities, or visual angles, display devices in non-contiguous configurations exist at different depths. Because the eye has to rapidly refocus when working at multiple depths, some ergonomics recommendations call for displays and documents to exist at a single depth (Ankrum, 1999). However, a study by Jaschinski-Kruza (1990) found that eyestrain was not increased when the user had to refocus their eyes at different depths. It should, however, be noted that a near sighted or far-sighted user has different abilities to see near objects or distant ones comfortably, so exceptions probably apply here (Chapanis & Scarpa, 1967).

Swanson and Couvillion (2001) report a study in which users had to divide their attention between different virtual depths. However, since they were primarily interested in comparing performance between various displays they made no effort to comment on the differences between working at multiple depths as compared to a single depth. In my work, I explicitly explore the effects of working on information at a single depth as compared to multiple depths in physical space.

6.3.1.2 Notifications

There has recently been a series of studies on the effects of notifications and other kinds of interruptions during everyday computing tasks (for a review, see McFarlane & Latorella, 2002). Most of these studies have shown the disruptive effects of notifications while multitasking (Czerwinski et al., 2000; Gillie & Broadbent, 1989; Kriefeldt & McCarthy, 1981; Maglio & Campbell, 2000). Gillie & Broadbent (1989) manipulated interruption length, similarity to the ongoing task, and the complexity of the interruption. They showed that even rehearsing the position of a target item in the main task does not protect a user from the disruptive effects of an interruption when trying to return to the target afterward. They also discovered that interruptions with similar content could be quite disruptive despite having an extremely short duration, replicating findings from earlier work by Kriefeldt and McCarthy (1981).

Other studies have examined the importance of spatial location of notifications, usually to determine the optimal display location for detection while minimizing disruption. For example, Hess, Detweiler, and Ellis (1999) showed that spatial locations were better than verbal labels, which were in turn better than visual-spatial icons, in supporting the temporary storage and retrieval of information. Their studies also showed that the number of notification updates was inversely related to memory performance for content.

Lim and Wogalter (2000) reported two studies that looked at the placement of static “banners” in a web browser window. In their first study, they examined banners in the extreme corners of the display and showed that recognition memory was significantly higher for banners placed in the top left or bottom right corners. Their second study showed that recognition performance was reliably higher for banners centrally located over those in the outer regions of the display. The authors argued that notifications could be made more salient by using this spatial location positioning. Unfortunately, the studies only utilized a single, 21" display, and did not explore larger or multiple display surfaces.

Bartram, Ware, and Calvert (2003) specifically explored notifications on larger displays using wider fields of view. The authors probed the perceptual properties of motion in an information-dense display with three experiments. They found that icons with simple motions are more effective than color and shape for notifications that must be delivered with low interruption. Based on these studies, they described several specific advantages and limitations of motion-based icons for larger displays. In addition, the authors varied the field of view affected during their detection tasks, making their guidelines and recommendations generalizable to larger display surfaces than the typical 17" to 21" monitors. However, the authors did not explore the effects of separation that hardware bezels and depth induce, and they focused only on design principles for notification detection. Here I examine multitasking performance while attending to and dismissing notifications across multiple displays.

6.3.2 Hypotheses

I ran a user study in order to systematically explore the effects of visual separation and physical discontinuities when distributing information across multiple displays while multitasking in a single user desktop scenario. Eccentricity effects suggest that the further two pieces of information are from each other in the visual field, the harder it is to divide attention between them. Thus,

Hypothesis a: Separating information by wider visual angles decreases task performance.

Even at equal visual angles, information divided by physical discontinuities such as monitor bezels or depth is harder to treat as a single unit and thus requires more cognitive resources for divided attention tasks. Hence,

Hypothesis b: Separating information by physical discontinuities decreases task performance.

6.3.3 Experiment

6.3.3.1 Participants

Twenty-four (12 female) users from the Greater Puget Sound area participated in the study. Users were intermediate to advanced Windows users with normal or corrected-to-normal eyesight. They ranged from 18 to 55 years of age (mean: 36.9). Users received software gratuity for their participation.




Figure 6.4. Experiment setup. I held visual angles constant between the Small and Large Display conditions. The primary display was always a small one.

6.3.3.2 Experiment and Setup

I used three displays, two NEC MultiSync FE1250 22" monitors and a Sanyo PLC-XP30 LCD projector. All displays ran at a resolution of 1024×768 and were calibrated to be of roughly equivalent brightness and contrast. The image on each monitor was 16" wide by 12.5" tall. The image projected on a wall-mounted screen was adjusted to be exactly 66" wide by 49.5" tall. One of the monitors was always the left display. I set up the second monitor and projection screen as the right display. When either of these displays was viewed from the user's seated position, the visual angle would be identical (see Figure 6.4). I assumed a comfortable viewing distance of 25" for the monitors. In order to get an image with identical visual angles, the large projection display was set up to be 103" away from the user. The center points of both displays were set to be at eye-height, about 60" above the ground. The position of the right monitor was carefully marked so that it could be moved in and out accurately for each condition.

I ran the study on a single 800 MHz Dell computer equipped with a dual-headed nVidia GeForce2 MX graphics card. I duplicated the output for the right display across the monitor and projector using an Inline IN3254 video splitter. Only one of the right displays was turned on at any given time. The user provided input with a standard keyboard and Microsoft IntelliMouse.

6.3.3.3 Tasks and Procedure

For this study, I created a compound test comprising a primary task performed in conjunction with a secondary and tertiary task. In the primary task, proofreading, users had to identify grammatical errors within a set of text articles. This task was chosen because it is not only visually but also cognitively demanding. I chose seven articles that appeared in the New York Times between January 1998 and December 2000. These articles were selected to be of similar readability and length. Flesch (1948) readability scores for the articles ranged from 46 to 54 (mean: 49.5), representing text at 11-12th grade reading level. Each article was at least 2000 words long.

I introduced errors into each article according to the following rules: (i) each sentence had at most one error, though some had none; (ii) errors were fairly evenly spaced throughout the article; (iii) errors included only subject-verb agreement, inconsistent verb tense, and word order (i.e. two words were flipped). These errors are similar to those introduced by Maglio and Campbell (2000) in their reading tasks. I instructed users to find as many errors as they could in the articles, marking each by double clicking on the word in question. They did not have to suggest corrections to the errors.

The secondary task is one I call notification detection. In this task, users had to detect and act upon visual changes outside the focal region of the primary task. This task is common, for example, in system notifications such as instant message arrival or print job completions, which are meant to keep users immediately aware of updated information. These notifications typically call for some form of user response. In my task, users had to detect a pop-up window modeled after the MSN instant messenger notification, and respond by hitting the space bar as quickly as possible.

Properly detecting a notification brought up the tertiary task, text comparison. Text comparison is representative of tasks in which the user must cross reference and compare content displayed in multiple locations on the displays. This is an important scenario since one of the benefits of having multiple displays is being able to view, compare and contrast more information simultaneously. In this task, a random set of 4 contiguous lines are selected from the text currently in view in the proofreading task. These lines are highlighted in the actual text as well as replicated in a dialog box which appears on the opposite display. The text in the dialog box is randomly chosen to be either a verbatim representation of the highlighted text or to have a single word order change. Users had to care-

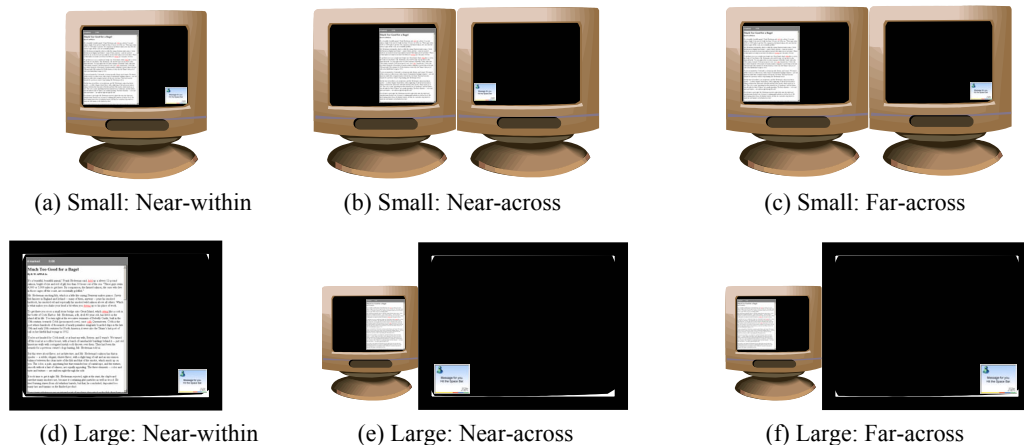


Figure 6.5. Display conditions used in the experiment.

fully compare the two sets of text and determine whether or not there was a change in the dialog box. They indicated their answer by clicking on one of two buttons, labeled ‘same’ or ‘different’ above the article. After doing this, they resumed proofreading.

In each trial, users were given 4 minutes for the proofreading task. Six notifications were randomly distributed with the constraint that they were at least 20 seconds apart. The clock that showed users how much time remained for proofreading was halted when a notification was detected. It was restarted after the user completed the text comparison task.

6.3.3.4 Design

I used a within subjects design. Each user performed 1 practice trial and 6 test trials, one in each of the 6 conditions, created using a 2 (Display Size: Small vs. Large) \times 3 (Distance: Near-within vs. Near-across vs. Far-across) design (see Figure 6.5).

The visual angle between the primary proofreading task and the secondary and tertiary tasks in the Near-within condition was kept exactly the same as in the Near-across condition (~ 27 degrees). The only difference between these two conditions was that the Near-within condition was completely contained within one display, whereas the Near-across condition was split across two, either having the monitor bezels or the bezels plus a depth discontinuity between the tasks. The Far-across condition was designed by keeping the primary task in the same position as in the Near-within condition and moving the secondary task as far away on the right display as possible (~ 55 degrees). The order of

conditions and articles used in the primary task were both counterbalanced using Latin Square designs.

Dependent measures included the number of errors correctly identified in the proof-reading task, the number of notifications correctly detected, the average reaction time of correctly detected notifications, the number of text comparisons correctly answered, and the average task time for these text comparisons. After the experiment, users filled out a preference survey, indicating the ease of performing the tasks in each of the conditions. The experiment took about an hour.

6.3.3.5 Results and Discussion

6.3.3.5.1 Overall MANOVA

I submitted the data to a 2 (Display Size: Small Display vs. Large Display) \times 3 (Distance: Near-within vs. Near-across vs. Far-across) repeated measures multivariate analysis of variance (MANOVA). Each dependent measure is covered separately in the results.

I observed no significant effects or interactions for either the average reaction time to detect a notification or the average reaction time for the text comparison task, at the $p=.05$ level (all effects were tested at this alpha level). Most users detected all the notifications and there were no significant effects with this measure.

For the number of correct text comparisons, I observed a significant interaction between Display Size and Distance, $F(2,46)=3.05$, $p=.05$. Post-hoc analyses showed that the




Figure 6.6. Though there were no significant differences on the Small Display for number of correct text comparisons (left), there was a significant difference between near-within and far across for the Large Display condition (right).




Figure 6.7. Though there were no significant differences on the Small Display for number of errors found (left), there was a significant difference between near-within and far across for the Large Display condition (right).

Near-within and the Far-across conditions were borderline significantly different, $p=.06$. The interaction reached significance because this difference between the Near-within and Far-across conditions was reliable for the Large Display (means: 5.167 and 4.625 respectively), though not the Small Display condition (means: 5.042 and 4.875), as can be seen in Figure 6.6. Although the result reaches statistical significance, the effect is fairly small.

For the number of correct errors found in the proofreading task, the interaction between size and distance reached borderline significance, $F(2,46)=2.6$, $p=.085$. Again this result was driven by a larger difference between the Near-within and Far-across conditions on the projection display (means: 7.875 and 7.000 respectively) but not the Small Display (means: 7.667 and 8.000), as seen in Figure 6.7. These effects are also relatively small.

These performance results ran counter to my initial hypotheses. I expected large, detrimental effects from separation of information in the visual field. I also expected detrimental effects from the physical discontinuity caused by the bezel and the separation in depth. For the time to detect notifications and for the text comparison times, I observed no effects of separation, bezel, or depth. In fact, I did not observe a significant main effect of visual separation in the performance data for any dependent measure. Instead, I observed a small but reliable interaction between display and the distance variable for the overall proofreading correct and text comparison correct measures. This interaction could

be best described as resulting from the differences between the Near-within and Far-across conditions being stronger for the Large Display condition.

6.3.3.5.2 *Satisfaction Data*

After the study, participants were asked which display configuration they preferred for performing the tasks involved. Surprisingly, 14 out of 24 participants stated that they preferred the smaller, 22" CRT for their primary task, significant by binomial test, $p=.006$. Nine participants preferred the larger, wall display for the primary task, and this was not significant. One participant stated “no preference” as their response to this question.

Participants were evenly split in terms of which configuration they preferred (same screen, split screen, or neither) for working on all experimental tasks. 10 preferred the tasks on the same screen, 11 preferred them on split screens, and 2 participants stated no preference.

This result is quite interesting, and converges nicely with some of the performance-based results I observed during the experiment. It appears that users are evenly split in how they would like their information presented around the bezel, and the deleterious effects appear to be much less important than I had hypothesized. The fact that about half the participants preferred to split their task across the bezels (even when distance to a larger, wall display is involved) is a fascinating one. I assert that the bezel might be playing some role that allows users to spatially address their information workspace in a way they perceive to be beneficial to the task, as asserted by Grudin (2001). However, this resulted neither in a reliable benefit nor detriment to task performance. Exploring this hypothesis more deeply remains as future work.

6.3.4 **Design Recommendations**

For the tasks chosen in this study, I saw significant performance differences between the Near-within and Far-across conditions, but only when information was split between the desktop monitor and the projection display. This indicates that, even at similar visual angles, placing information further in the periphery on displays that are separated in depth is more detrimental to performance than the corresponding position at similar depths. However, it should be noted that for my tasks, effects seen were relatively small (about a 10% performance decrement), and designers, aware of the differences present, can weigh the importance of the information to be displayed with this trade-off in mind.

Interestingly, I saw no effects of physical discontinuities, introduced either by monitor bezels or by the depth difference between the Small Display and projection display. This was surprising, but implies that designers might have more freedom when splitting information across boundaries than I had anticipated. I do not doubt that there are tasks which will be hurt by splitting information across physical discontinuities, but my set of tasks (proofreading and monitoring) do not seem to fall heavily into that category.

6.3.5 Summary and Future Work

In this section, I have reported a study examining the effects of visual separation and physical discontinuities when distributing information across multiple displays. Study tasks were chosen to be representative of tasks carried out by information workers while multitasking so as to increase the generalizability of the results to future display systems and user interface designs for single user desktop scenarios. The study demonstrated that there is a reliable, though relatively small, detrimental effect on performance from separating information within the visual field when it is further separated by depth. Also, counter to my hypotheses, physical discontinuities introduced by bezels as well as by differences in depth alone do not seem to have an effect on performance on the set of tasks I have chosen. I have presented design recommendations that follow from these results.

I would like to extend this work in several directions. I would like to add further ecological validity by introducing unrelated notification content that serves as extra distraction. In the current study, I displayed only information that was relevant to the tasks the user was performing. Previous research has shown that this should make visual search and detection tasks harder (Czerwinski et al., 2000; Gillie & Broadbent, 1989), but I do not know the effects of my manipulations in this situation.

Also, more work needs to be done to explore scenarios that involve collaboration and interruption, as well as different tasks within the same experimental framework. For example, I could use a monitoring task, in which users have to simultaneously watch and act upon multiple objects while communicating and sharing information with other colleagues. Alternatively, I could extend this work to tasks in which depth cues or continuity of the information is important, such as in certain 3D environments. Results from this work have critical implications both on the design of large display workplaces as well as on software and applications operating in these new display configurations.

"I reach a conclusion whenever I am tired of thinking"

Anonymous

*"What we call the beginning is often the end.
And to make an end is to make a beginning.
The end is where we start from."*

T.S. Elliot

Chapter 7

Conclusion

7.1 Summary of Work and Specific Contributions

With the diversification of computing devices, the continued increase in processing power, the miniaturization of components, and the widespread explosion of wireless connectivity, we have seen enormous interest in research areas such as ubiquitous computing (Weiser, 1991), pervasive computing (Husemann, 2000), and personal technologies (Frolich, Thomas, Hawley, & Hirade, 1997). All these fields aim to shift the computing paradigm to one that is more closely embedded within people's lives. In fact, there is a growing search for useful paradigms, culminating in research on tangible user interfaces (Ishii & Ullmer, 1997), ambient media (Wisneski et al., 1998), information appliances (Norman, 1998), context awareness (Schmidt, 2000), invisible computing (Weiser & Brown, 1996), the disappearing computer (European Commission, 2000), and ensemble computing (Thomas & Gellersen, 2001).

The work I have presented in this dissertation has been motivated by a goal very much in line with these emerging research areas. In fact, my broad agenda has been defined by the desire to craft computing environments that more fully support human thought and action. In doing this, I have realized that cognition and action are necessarily situated within the larger context of the physical environment. Hence, I believe that creating more effective computing systems requires us to understand how we can best leverage some of the physical properties of the environment. In my work, I have chosen to use

physically large displays to motivate thought around these issues. While the focus of this work has centered on physically large displays, it should be noted that this is only one piece of a much larger puzzle. Hence the immediate contributions of this dissertation should be considered within the context of the growing movement to think of computers not as devices, but as environments.

Work reported in this dissertation can be broken into three basic components. Each component is a fairly distinct set of work and each contributes uniquely to the overarching goal. In the following sections, I summarize the contributions of this work.

7.1.1 Contributions to Theoretical Understanding of Cognitive Benefits

While researchers have carefully explored display characteristics such as field of view, resolution, brightness, contrast, and color, little has been done to examine the effects of physical display size. This is partially due to the fact that display technologies used in traditional computing environments have seldom extended significantly beyond the form factor of a desktop monitor. Furthermore, even when displays have taken on radically different form factors within complex display environments, researchers have typically opted to study them by treating them as an integrated whole rather than by decomposing and examining specific characteristics. While this has been a relatively productive approach thus far, I believe that we should augment this with work aimed at understanding the specific effects of individual display characteristics. By doing this, we can begin to build a more general theory that is useful not only in understanding a wide range of current display technologies, but also in developing new ones. Aside from technological benefits, this work also provides insight into interesting psychophysical phenomena that I hope will extend the way we think about human cognition within our physical and digital environments.

In Chapter 3, I described a series of experiments that provide insight into the effects of physical size. I used both standard psychology tests as well as more ecologically valid tasks to show that physical display size, independent of other factors, elicits cognitive and behavioral responses that can affect task performance on spatial tasks. Specifically, the contributions of this work include:

- The identification of physical display size as an important display characteristic that must be considered in designing our future display systems.

- Experimental results suggesting that physical display size biases users into adopting different strategies when they perform spatial orientation and visualization tasks. In fact, small displays seem to bias users into exocentric strategies, and large displays seem to immerse users more within virtual environments and bias users into egocentric strategies.
- Experimental results showing that egocentric strategies only aid performance on tasks which benefit from having users imagine their bodies within the problem space.
- Experimental results suggesting that the benefits of large displays are independent of other display factors that could aid task performance, such as interactivity and additional cues present with the virtual environments.
- Application of the current understanding, showing that large displays can be used to improve performance on 3D navigation as well as mental map formation and memory tasks, both in controlled as well as ecologically valid virtual environments.

7.1.2 Contributions to Tools Leveraging Social Affordances

Many researchers have articulated the utility of large displays for collaborative work. Large displays have the potential to provide public surfaces that offer greater visibility to more people. This is associated with reducing both technical as well as social barriers and is usually assumed to facilitate collaboration. While there have been many tools built for co-located users to work on shared displays, few of these have explicitly addressed the problem of scarce screen space when multiple people try to work on a single display. In fact, scarce screen space is a largely unsolved problem that exists even when a single user works on a traditional desktop display system. I believe that the ideal solution would provide users with the flexibility to specify the information that they want to view, and then to spatially arrange this information to optimally support the task at hand. Additionally, such a solution would also allow users to share information between devices so that they can utilize all the display resources within the environment.

In Chapter 4, I described WinCuts, an interaction technique I developed to explore the viability of such a solution. Although initially motivated by thinking about collaboration on physically large displays, the technique serves as a much more general purpose tool for managing information in general and working within rich display environments. The specific contributions of this work include:

- An interaction technique called WinCuts that allows users to manipulate arbitrary regions of information contained within existing application windows. WinCuts promotes a finer granularity of spatial screen space management, as well as sharing chunks of information between machines.
- A software artifact that shows one method of implementing the WinCuts system, but that also suggests the need for a more robust implementation of the model-view-controller architecture within our interfaces.
- The exploration of a fairly wide range of scenarios, with and without large displays, in which WinCuts may be used, suggesting the generality of the technique.

Another social phenomenon that emerges around the use of large displays is the intrinsic loss of privacy for content. Although many researchers have articulated the effects that exist, few have been able to quantify them because of the lack of a method to do so. In Chapter 5, I explored some of the social affordances of physically large displays, and contribute:

- An experimental method of measuring whether or not a person has seen and cognitively processed a particular piece of text using a novel application of an implicit memory priming paradigm.
- Experimental results showing that people are socially more likely to read content when it is shown on a physically large display, even when visual angles and legibility are equivalent to smaller displays.

One implication of these findings is that private information and actions on large displays are intrinsically treated as being public. Since current interfaces are not designed with this in mind, it is sometimes inconvenient to perform private tasks such as entering passwords on large touch screen displays. In the second half of Chapter 5, I described an interface designed with a sensitivity to the physical affordances and social effects of a specific device on which it will run. Here I contribute:

- An interaction technique called the Spy-Resistant Keyboard that allows users to type private text on a publicly observable virtual keyboard without revealing the text typed.

- Experimental results describing the costs and benefits of using such a system. These results suggest that even though security is drastically improved, ease of use requires substantial improvement.
- Design principles for building other such interfaces that protect typists from casual observers.

7.1.3 Contributions to Pragmatics

Finally, in working within the Display Garden, I have encountered and begun to explore some of the pragmatic issues surrounding the integration of large displays into our workspaces. While this work did not form a significant portion of my agenda, I believe that it is important as it highlights two of the more obvious issues faced within environments of this sort.

In Chapter 6, I described two pieces of work. Here I contribute:

- A method using unobtrusive infra-red light and computer vision for sensing the presence of a person in front of a projector and eliminating blinding light from being cast on that person.
- Experimental results showing the effects of visual separation and physical discontinuities that exist in multiple display environments with large displays. In fact, results show small, but detrimental effects on performance when information is separated within the visual field, but only when it is further separated by depth. Also, bezels and depth alone did not seem to have any effect on the set of tasks tested.

7.2 Future Work

I have presented specific pieces of future work within the respective chapters. I will close my dissertation with a brief overview of my longer term agenda. This agenda is driven by the desire to move away from building standalone computing devices connected to each other mainly as an afterthought, to crafting much richer computing environments consisting of multiple integrated components. Ideally, these would be environments in which multiple users could operate, each utilizing the input, display, and processing resources that are most helpful in completing the tasks at hand.

I believe that there are many challenges to attaining this goal. The largest challenge will be defining the metaphor that will enable users to operate effectively in such an environment. The metaphor must capture the fluid nature of the system, both with respect to transient ad hoc configuration of devices, but also to interaction techniques that work across the entire computing landscape in order to support task performance. The metaphor must also create a sense of transparency across the various component technologies, so that users do not concern themselves so much with coordinating multiple devices as they do completing their tasks with the help of the environment. It is this metaphor that will allow the new model to gain traction within the computing community, but more importantly it is this metaphor that will define the new capabilities that such environments will offer users.

One useful step that I believe must be taken is reconsideration of the way we architect and engineer current devices and hardware systems. Since we have traditionally assumed that each device is a standalone computing entity, we have not placed enough emphasis on proper modular construction to best support software applications as well as functional integration. For example, it is currently very difficult for a user to walk into a random office and set up their laptop so that it uses all possible input devices and display space in that environment. I believe that there is huge potential in constructing machines such that input and display devices as well as processing capabilities are better separated and easily addressable by other such units. At a high level, this would allow users to use any input device to control any machine which can then show output on any display device in the environment at any time. As an additional pragmatic issue, I believe that in this reconsideration of modular hardware, rendering capabilities must be more tightly associated with their specific display devices so that displays are no longer just dumb devices that do nothing but draw pixels.

In the longer term, I also hope to tackle other issues involved in building systems that allow users to seamlessly take advantage of screen space, specialized input devices, processing capabilities, portability, or other affordances of all their devices. These issues include building automatic discovery mechanisms that support dynamic configuration of devices, architecting efficient communication between devices, exposing the structure of the environment, designing interfaces that let users configure devices, and building brand new applications that take full advantage of operating across multiple devices.

The end goal, hopefully extending well beyond the scope of this dissertation, is to re-orient our role as human-computer interaction researchers so that it includes a much broader perspective of cognition and technology. I believe that success as a field will necessarily entail shifting some attention to physical properties of information displays and input devices in the environment in order to acquire a better understanding of our interaction with computers in the digital and physical realms. With this understanding, we can then craft the tools and environments that best support human thought and action.

Appendix A: Selected Experiment Materials

A.1 Pre-Test Questionnaire (used across all experiments)

User Number (*filled by experimenter*): _____

Date/Time: _____

Age: _____

Gender (*circle*): Male Female

Profession: _____

Highest education level (or current level if still in school): _____

Major: _____

Handedness (*circle*): Right handed Left handed Ambidextrous

Do you wear (*circle all applicable*): Contact Lenses Glasses

Which are you wearing today (*circle*): Contact Lenses Glasses

Average estimated computer use per week: _____ hours

Rank order the frequency of use for the following displays (leave any that you don't use blank):

_____ CRT Monitor (19" or smaller)



_____ CRT Monitor (greater than 19")

_____ LCD Monitor (19" or smaller)



_____ LCD Monitor (greater than 19")

_____ Projector



_____ Other(s), please specify: _____

A.2 Post-Test Questionnaire (variants used across experiments)

User Number (filled by experimenter): _____

For the Spatial Task (determining rotation of the boats)

1. For each of the following factors (circle preference):

Ease of seeing:

Strongly prefer small display	Prefer small display	Don't care	Prefer large display	Strongly prefer large display
----------------------------------	-------------------------	------------	-------------------------	----------------------------------

Confidence in rotation task:

Strongly prefer small display	Prefer small display	Don't care	Prefer large display	Strongly prefer large display
----------------------------------	-------------------------	------------	-------------------------	----------------------------------

Overall preference:

Strongly prefer small display	Prefer small display	Don't care	Prefer large display	Strongly prefer large display
----------------------------------	-------------------------	------------	-------------------------	----------------------------------

2. Comments:

For the Reading Task

3. For each of the following factors (circle):

I found it comfortable reading the text:

Small Display

Small text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Medium text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Large text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree

Large Display

Small text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Medium text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Large text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree

I think that I read passages quickly:

Small Display					
Small text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Medium text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Large text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Large Display					
Small text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Medium text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Large text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree

I think that I did well on the questions:

Small Display					
Small text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Medium text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Large text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Large Display					
Small text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Medium text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree
Large text	Strongly Disagree	Disagree	Indifferent	Agree	Strongly Agree

4. Overall, I preferred the (circle):

Small display Large Display No preference

5. Overall, I preferred the (circle):

Small text Medium text Large text No preference

6. Comments:

In General

7. What did you like about the small display?
8. What did you NOT like about the small display?
9. What did you like about the large display?
10. What did you NOT like about the large display?
11. Why did you prefer the text size that you selected?
12. Other comments:

A.3 Instructions for Guilford-Zimmerman Spatial Orientation Task (Original Version)

This is a test of your ability to see changes in direction and position. In each item you are to note how the position of the boat has changed in the second picture in relation to the original position in the first picture.


Here is Sample Item 1

These bars represent the prow of the boat.

This is the correct answer. It shows that the prow of the boat is now below the aiming point.

(If the prow had risen, instead of dropped, the correct answer would have been 3, instead of 4)

There are five possible answers to the item.



This is the prow (front end) of a motor boat that you are standing on.

This is the aiming point. It is the exact spot you would see on the painted backdrop if you sighted right over the point of the prow.


This is the same aiming point shown above. Note that the prow has dropped below it.

To work each item: First, look at the top picture and see where the tip of the motor boat is pointing (the aiming point). Second, look at the bottom picture and note the CHANGE in where the boat is pointing. Third, select the answer (1, 2, 3, 4, or 5) that best shows that change with the input device and press enter to confirm your selection.

Here is Sample Item 2

This also shows that the prow of the boat is to the right of the aiming point. So, it is the correct answer.

(If the boat had moved to the left, instead of to the right, the correct answer would have been 1)




This is the aiming point.

This is the same aiming point. The boat is now pointing to the right of it.


Here is Sample Item 3

This is the correct answer. It shows that the motor boat changed its slant to the left, but is still heading toward the aiming point.



Here the motor boat is slanted slightly to the right. (note that the horizon appears to slant in the opposite direction)


Here the boat has changed its slant toward the left. (To become level, the boat slanted back toward the right)



Imagine that these pictures were taken with a motion picture camera. The camera is fastened rigidly to the boat so that it bobs up and down and turns and slants with the boat. Thus, when the boat tips or slants to the left (as in the lower sample, sample item 3), the scene through the camera viewfinder looks slanted like this.

Here is Sample Item 4

4 is the correct answer. It shows that the boat has changed its heading both downward and to the right; also that it changed its slant towards the right.



The prow of the boat has moved downward and toward the right. Also, it changed its slant toward the right.

Now do Practice Items 1, 2, and 3 (these questions will not be evaluated).

The aiming point is not marked in the test items. You must see changes in the boat's position without the aid of the dots.

To select your answer, hit one of the number keys (1, 2, 3, 4, or 5) on the keypad. You may change your answer selection as many times as you like. To confirm, hit the "Enter" key.

To review:

First: Look at the top picture. See where the motor boat is heading.

Second: Look at the bottom picture. Note the change in the boat's heading.

Third: Use the number keys to select the answer that shows the same change (in reference to the aiming point before the change). Hit "Enter" to continue.

Good luck...

A.4 Instructions for Guilford-Zimmerman Spatial Orientation Task (Egocentric Version)

This is a test of your ability to see changes in direction and position. Imagine that you are a film director in a studio set. You are standing on top of a boat that is on the movie set. The crew is moving the boat as you are on the boat. Two pictures are taken, one before the boat moves and one after. In each item you are to note how the position of the tip of the boat has changed in relation to the painted backdrop.


Here is Sample Item 1

These bars represent the prow of the boat on the movie set.

This is the correct answer. It shows that the prow of the boat is now below the aiming point.

(If the prow had risen, instead of dropped, the correct answer would have been 3, instead of 4)

There are five possible answers to the item.



This is the prow (front end) of a motor boat that you are standing on.

This is the aiming point. It is the exact spot you would see on the painted backdrop if you sighted right over the point of the prow.


This is the same aiming point shown above. Note that the prow has dropped below it.

To work each item: First, look at the top picture and see where the tip of the motor boat is pointing (the aiming point). Second, look at the bottom picture and note the CHANGE in where the boat is pointing. Third, select the answer (1, 2, 3, 4, or 5) that best shows that change with the input device and press enter to confirm your selection.

Here is Sample Item 2

This also shows that the prow of the boat is to the right of the aiming point. So, it is the correct answer.

(If the boat had moved to the left, instead of to the right, the correct answer would have been 1)




This is the aiming point.

This is the same aiming point. The boat is now pointing to the right of it.


Here is Sample Item 3

This is the correct answer. It shows that the motor boat changed its slant to the left, but is still heading toward the aiming point.



Here the motor boat is slanted slightly to the right. (note that the horizon appears to slant in the opposite direction)


Here the boat has changed its slant toward the left. (To become level, the boat slanted back toward the right)



Imagine that these pictures were taken with a motion picture camera. The camera is fastened rigidly to the boat so that it bobs up and down and turns and slants with the boat. Thus, when the boat tips or slants to the left (as in the lower sample, sample item 3), the scene through the camera viewfinder looks slanted like this.

Here is Sample Item 4

4 is the correct answer. It shows that the boat has changed its heading both downward and to the right; also that it changed its slant towards the right.



The prow of the boat has moved downward and toward the right. Also, it changed its slant toward the right.

Now do Practice Items 1, 2, and 3 (these questions will not be evaluated).

The aiming point is not marked in the test items. You must see changes in the boat's position without the aid of the dots.

To select your answer, hit one of the number keys (1, 2, 3, 4, or 5) on the keypad. You may change your answer selection as many times as you like. To confirm, hit the "Enter" key.

To review:

First: Look at the top picture. See where the motor boat is heading.

Second: Look at the bottom picture. Note the change in the boat's heading.

Third: Use the number keys to select the answer that shows the same change (in reference to the aiming point before the change). Hit "Enter" to continue.

Good luck...

A.5 Instructions for Guilford-Zimmerman Spatial Orientation Task (Exocentric Version)

This is a test of your ability to see changes in direction and position. Imagine that you are a film director in a studio set. You are standing on top of a boat firmly attached to the floor of the movie set. The crew is moving a painted backdrop on the set. Two pictures are taken, one before the painted backdrop is moved and one after. In each item you are to note how the position of the painted set backdrop has changed relative to the tip of the boat on the movie set.


Here is Sample Item 1

There are five possible answers to the item.

These bars represent the prow of the boat on the movie set.

This is the correct answer. It shows that the prow of the boat is now below the aiming point.

(If the prow had risen, instead of dropped, the correct answer would have been 3, instead of 4)




To work each item: First, look at the top picture and see where the tip of the motor boat is pointing (the aiming point). Second, look at the bottom picture and note the CHANGE in the aiming point. Third, select the answer (1, 2, 3, 4, or 5) that best shows that change with the input device and press enter to confirm your selection.

Here is Sample Item 2


This also shows that the aiming point is to the left of the prow of the boat. So, it is the correct answer.

(If the aiming point had moved to the right, instead of to the right, the correct answer would have been 1)



Here is Sample Item 3

This is the correct answer. It shows that the painted backdrop changed its slant to the right, but did not move from the tip of the boat. (note: the answers have been slightly rotated. Also, imagine that the aiming point has a certain orientation and remember the line represents the boat on the stage)



Here the painted backdrop is rotated slightly to the left.


Here the painted backdrop has rotated towards the right (and the boat appears to be headed left)



Imagine that these pictures were taken with a motion picture camera. The camera is fastened rigidly to the boat. The image is of the painted backdrop and it can be moved up or down, left or right, and rotated around a point that is on the tip of the boat on the movie set. Thus, when the painted backdrop rotates to the right (as in the lower sample, sample item 3), the scene through the camera view finder looks like this.

Here is Sample Item 4

4 is the correct answer. It shows that the aiming point has changed its direction both upward and to the left; also that it changed its slant to the left.



The aiming point has moved upward and toward the left. Also, it changed its slant toward the left.

Now do Practice Items 1, 2, and 3 (these questions will not be evaluated).

The aiming point is not marked in the test items. You must see changes in the boat's position without the aid of the dots.

To select your answer, hit one of the number keys (1, 2, 3, 4, or 5) on the keypad. You may change your answer selection as many times as you like. To confirm, hit the "Enter" key.

To review:


First: Look at the top picture. See where the aiming point is pointing.

Second: Look at the bottom picture. Note the change in the aiming point's (and the painted backdrop) direction and heading.


Third: Use the number keys to select the answer that shows the same change (in reference to the aiming point before the change). Hit "Enter" to continue.

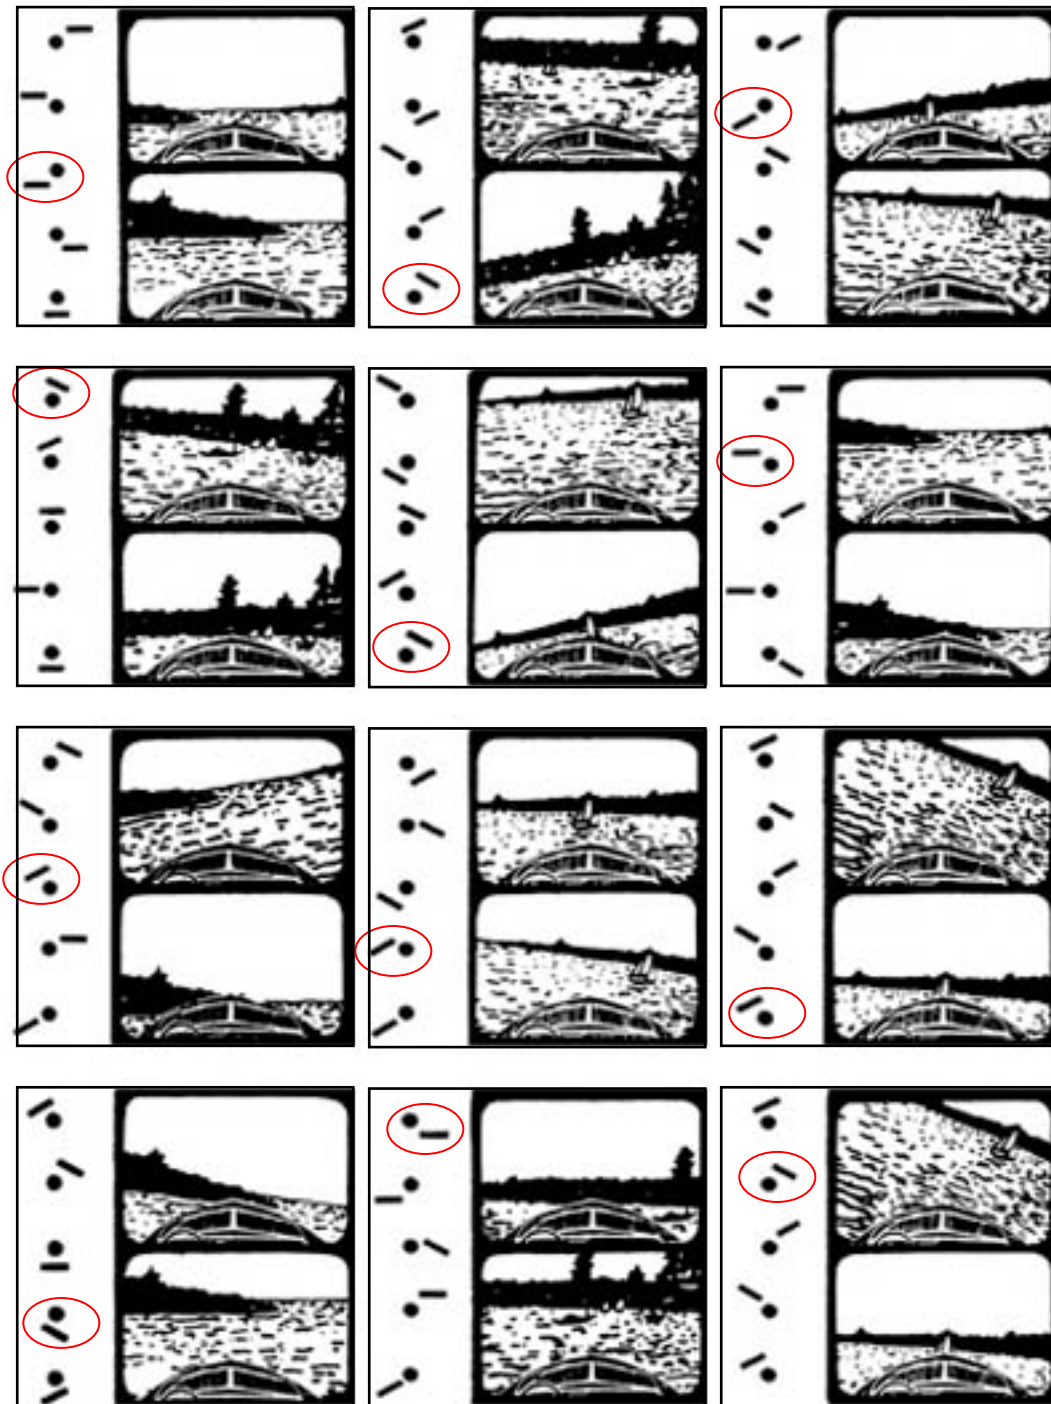
Good luck...


A.6 Stimuli for Guilford-Zimmerman Spatial Orientation Task















A.7 Instructions for Card Test

This is a test of your ability to see differences in figures. Look at the 5 triangle-shaped cards.




All of these drawings are of the same card, which has been slid around into different positions on the page. Now look at the 2 cards below: These two cards are not alike. The first cannot be made to look like the second by sliding it around on the page. It would have to be flipped over or made differently.




Each question consists of 8 pairs of cards presented one after the other. For each of these 8 pairs, the left card will remain the same. Only the right card will change. For each pair, you are to decide whether each of the cards on the right is the same or different from the card at the left.









Practice on the following pairs:










The pairs on the top row should have been marked “Same,” and the pairs on the bottom should have been “Different.” Make sure you understand why before proceeding.

You will be paid a performance bonus based on speed as well as accuracy. Therefore, work as quickly as you can without sacrificing accuracy. You will have 3 minutes for each of the two parts of this test. Each part has 10 questions, each with 8 pairs, for a total of 80 pairs. When you are done, the experimenter will start the next part. Good luck...


A.8 Stimuli for Card Test






A.9 Instructions for Cube Test

Wooden blocks such as children play with are often cubical with a different letter, number, or symbol on each of the six faces (top, bottom, four sides). Each problem in this test consists of drawings of pairs of cubes or blocks of this kind. Remember, there is a different design, number, or letter on each face of a given block or cube. Compare the two cubes in each pair below.



Example 1: Different




Example 2: Same

The first pair is marked “Different” because they must be drawings of different cubes. If the left cube is turned so that the A is upright and facing you, the N would be to the left of the A and hidden, not to the right of the A as is shown on the right hand member of the pair. Thus, the drawings must be of different cubes. The second pair is marked “Same” because they could be drawings of the same cube. That is, if the A is turned on its side the X becomes hidden, the B is now on top, and the C (which was hidden) now appears. Thus the two drawings could be of the same cube.

No letter, number, or symbol appears on more than one face of a given cube. Except for that, any letter, number, or symbol can be on the hidden faces of a cube.

Practice on the three examples below:




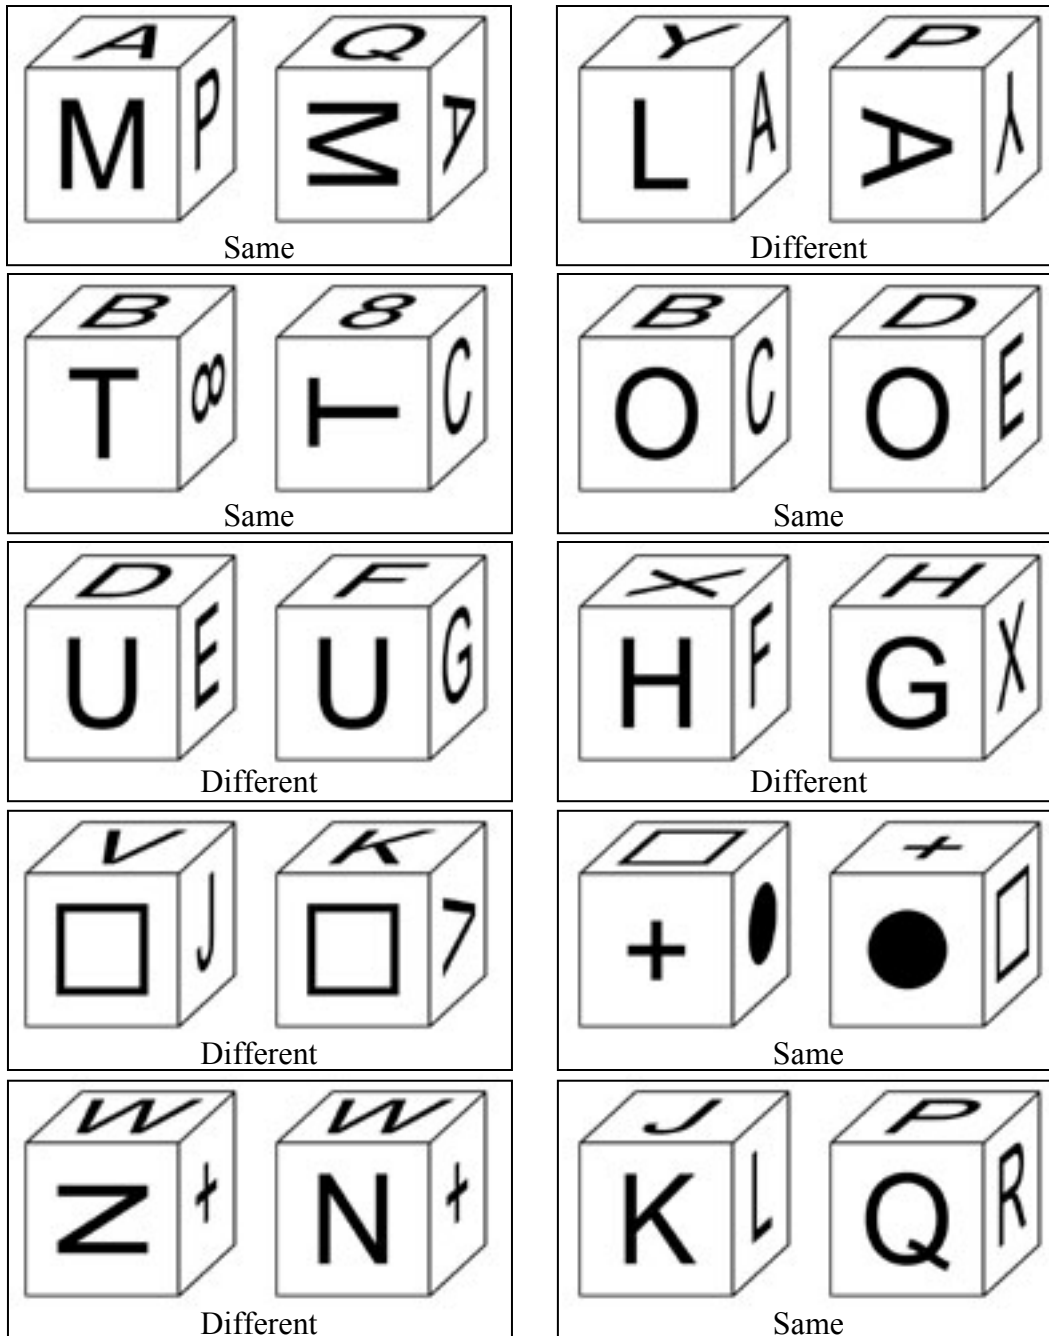
The first pair immediately above should be marked “Different” because the X cannot be at the peak of the A on the left hand drawing and at the base of the A on the right hand drawing. The second pair is “Different” because P has its side next to G on the left hand cube but its top next to G on the right hand cube. The blocks in the third pair are the “Same”, the J and K are just turned on their side, moving the O to the top. Make sure you understand this before proceeding.











You will be paid a performance bonus based on speed as well as accuracy. Therefore, work as quickly as you can without sacrificing accuracy.











You will have 3 minutes for each of the two parts of this test. Each part has 21 questions. When you are done with each part, the experimenter will start the next part. Good luck...








A.10 Stimuli for Cube Test













A.11 Instructions for Shepard-Metzler Test

In this test, you will see pairs of objects in space. Each object consists of 10 solid cubes attached face-to-face to form a rigid arm-like structure. Each object is also rotated to varying degrees.

Your task is to determine if the two objects seen in each pair are the “Same” object, or if they are “Different”. For example,







Example 1: Same



Example 2: Different

The first pair should be marked “Same” since the pictures can be generated by rotating the same object in space. The second pair, however, should be marked “Different” since there is no rotation that can make the two objects line up perfectly. In fact, in example 2, the objects are not the same physical objects; they are mirror images of each other.

Practice on the four examples below:



For the same reasons as before, you should have marked the pairs on the left “Same” and the pairs on the right “Different”. Make sure you understand why before proceeding.

You will be paid a performance bonus based on speed as well as accuracy. Therefore, work as quickly as you can without sacrificing accuracy.

You will do 2 sections of this test, each with 60 questions. There is no time limit for this test. When you are done with each part, the experimenter will start the next part. Good luck...

A.12 Stimuli for Shepard-Metzler Test

The complete set of 70 line drawings consists of 10 different objects (2 subsets of 5; namely 5 objects and their corresponding mirror images projections) in 7 positions of rotation about a vertical axis. Each drawing is coded in the following manner: the upper case letters, “A”, “B”, “C”, “D”, and “E” denote the 5 different objects in each subset; a lower case “p” or “n” refers to a “positive” or “negative” (i.e., mirror image) projection; and the numbers designate the perspective view. (Thus Ap6 and An6 are mirror images of one another.)

The numbers of the perspective views for each object, as well as the angular differences between them, are indicated below.

Object	Angular differences (in degrees) between projections							
	20°	80°	80°	20°	40°	80°	40°	
A:	1	2	6	10	11	13	17	1
B:	14	15	1	5	6	8	12	14
C:	5	6	10	14	15	17	3	5
D:	10	11	15	1	2	4	8	10
E:	5	6	10	14	15	17	3	5

For example, a pair of drawings consisting of Ap1 and Ap6 would constitute a SAME pair (i.e., the objects can be rotated into congruence), differing in orientation by 100°, whereas Ap1 and An6 would represent a DIFFERENT pair) i.e., the objects cannot be brought into congruence by any rotation).

The 7 perspective views for each object permit the construction of at least 2 unique pairs at each angular difference in orientation from 0° to 180°, in 20° steps.

NOTE: The figure labels (e.g., Ap1, An6, etc.) should be removed or concealed during use in experimentation.




An1



Ap1



An2



Ap2



An6



Ap6



An10



Ap10



An11



Ap11



An13



Ap13



An17




Ap17



Bn1



Bp1



Bn5



Bp5



Bn6



Bp6



Bn8



Bp8



Bn12



Bp12



Bn14



Bp14



Bn15



Bp15



Cn3



Cp3



Cn5



Cp5



Cn6



Cp6



Cn10



Cp10



Cn14



Cp14



Cn15



Cp15




Cn17



Cp17



Dn1



Dp1



Dn2



Dp2




Dn4



Dp4



Dn8



Dp8



Dn10



Dp10



Dn11



Dp11



Dn15




Dp15



En3



Ep3



En5




Ep5




En6



Ep6



En10




Ep10



En14



Ep14



En15



Ep15



En17



Ep17


A.13 Instructions for Path Integration Test

Practice Task 1


This part of the experiment will familiarize you with the test setup and input device, as well as provide practice on representative tasks that you will perform during the actual test. We will not analyze data from this part of the experiment so feel free to play around and get comfortable. Also, ask questions as you progress through this phase.

For the experiment, you will be moving through a simple 3D virtual environment using a joystick. Pushing the joystick forward/backward will move you forward/backward; pushing the joystick to the left/right will turn you left/right. You cannot move and turn at the same time. Hitting the trigger (on the back of the joystick) starts and ends tasks.

There will be three red poles for each question in the practice task. The poles are set up to be the corners of an imaginary triangle (below figures). You will always move in a counterclockwise direction, starting at the southernmost pole. This pole will not be visible most of the time. You will be led along the first two legs of the triangle by following the other two poles, which will show up one at a time. Only one of the poles will be visible at each point in time. After traveling the first two legs, your task will be to turn and move to the point from which you started. You will have no visible poles to guide you in this final stage.



With the aid of the poles, travel the first two legs of the triangle.



Without the aid of poles, turn back to the start point and then move to it.



Detailed Instructions



Step 1: Click trigger to begin the question

Step 2: View overview map of the environment and triangle with all poles visible.
The blue sphere is where you will start (above figure)

Step 3: Click trigger to begin the actual task



Step 4: Push joystick forward to move straight ahead to the first pole

Step 5: Push joystick right to face the second pole (hint: this will always be on your right). It will light up a bright red when you're actually facing it (above figures).

Step 6: Push joystick forward to move to the second pole



Step 7: Push joystick right to turn to the unmarked start point

Step 8: Push joystick forward to move to the start point

Step 9: Click on the trigger when you think you're back at the start point. You will see another overview map showing where you ended up (above figure).

Repeat (there will be 6 questions)

Practice Task 2

As with the last task, we will not analyze data from this part of the experiment. Feel free to ask questions as you progress through this phase.

This practice task will be the same as the last except that you will **not** control movement along the first two legs. Instead, the computer will control this movement. You will still have to turn and move to the start point.

Note: You will **not** get overview maps in the actual task, so do not get too reliant on them. The maps are purely to help you get used to the task.

Step 1: Click trigger to begin the question

Step 2: View overview map of the environment and triangle with all poles visible. The blue sphere is where you will start

Step 3: Click trigger to begin the actual task

~~Step 4: Push joystick forward to move straight ahead to the first pole~~

~~Step 5: Push joystick right to face the second pole (hint: this will always be on your right). It will light up a bright red when you're actually facing it.~~

~~Step 6: Push joystick forward to move to the second pole~~

Step 4-6: Watch movement along the first two legs

Step 7: Push joystick right to turn to the unmarked start point

Step 8: Push joystick forward to move to the start point

Step 9: Click on the trigger when you think you're back at the start point. You will see another overview map showing where you ended up.

Repeat (there will be 6 questions)

Test Instructions

Again, you will be led to walk along 2 legs of an imaginary triangle and will then be asked to complete the triangle and return to the start point without any guidance. You will perform 6 questions in each of 4 different conditions. Everything will be the same as the practice tasks except that (a) you will not get an overview map either before or after each question (in fact, you will get no feedback), (b) you will sketch the triangle in between questions. The experimenter will briefly remind you of each condition before you begin.

Your performance on part of the experiment will be analyzed. You should perform all tasks as quickly and accurately as you can. Sketch the triangles between questions. Your bonus compensation will be calculated based on your performance (speed and accuracy) on one of these 4 conditions. The actual condition that will be used will be randomly selected at the end of the experiment.

As a reminder:

~~Step 1: Click trigger to begin the question~~

~~Step 2: View overview map of the environment and triangle with all poles visible.~~

~~The blue sphere is where you will start~~

Step 3: Click trigger to begin the actual task

Step 4: Push joystick forward to move straight ahead to the first pole

Step 5: Push joystick right to face the second pole (hint: this will always be on your right). It will light up a bright red when you're actually facing it.

Step 6: Push joystick forward to move to the second pole

or

Step 4-6: Watch movement along first two legs

Step 7: Push joystick right to turn to the unmarked start point

Step 8: Push joystick forward to move to the start point

~~Step 9: Click on the trigger when you think you're back at the start point. You will see another overview map showing where you ended up.~~

Repeat (there will be 6 questions per condition)

A.14 Instructions for Map Formation and Memory Test (Variants used for Controlled and Ecologically Valid versions)

Overview

Thank you for taking time out and coming in to help us. In this study, we are looking at how well people are able to navigate within 3D virtual environments. Your results and comments will help guide our design of future hardware and software systems.

Today, you will be performing a series of simple navigation tasks in several different environments. In most cases, you will be using the Microsoft joystick on the table in front of you. We are currently interested in how well people use the joystick to move around, so we will only use the main control stick. None of the dials and buttons will be used.

Control

Moving through the environment is simple. To move forward/backward, push the stick forward/backward. To turn left/right, push the stick left/right. Once you get good at this you can move and turn at the same time. For example, pushing the stick forward and right at the same time will turn you to the right as you move forward.

Task

Each environment contains 4 large dice, labeled 1 through 4. There are also random walls that form obstacles throughout the environment. For each environment, you will be allowed to wander around for 4 minutes. Your task during this period is to learn the positions of the dice and the layout of the environment as well as you can.

After the 4 minutes, you will be placed in the center of the environment and asked to find and walk through a particular cube. You'll repeat this 8 times, twice for each cube. We will be measuring (a) how quickly you get to the cube; (b) how short a path it takes you to get there. So work quickly and try to find the shortest path to the cube.

Finally, at the end of each environment, you will have to select the overview map that you think best represents the environment (from a set of 4).

The fun begins...

The first environment is purely for you to practice the controls and the task. We will not be looking at the result from this environment. So feel free to play around and get used to it all. After this, we'll stop for you to clarify anything that is confusing. We'll do 4 test environments (the experimenter will give you additional instructions for some of these), a quick questionnaire, and then we're done!

Good luck...

Bibliography

- Ackerman, M.S. (2000). The intellectual challenge of CSCW: The gap between social requirements and technical feasibility. *Human-Computer Interaction*, 15, 179-203.
- Adams, A., Sasse, M.A., Lunt, P. (1997). Making passwords secure and usable. *Proceedings of HCI on People and Computers XII*, 1-19.
- Alfano, P.L., Michel, G.F. (1990). Restricting the field of view: Perceptual and performance effects. *Perceptual and Motor Skills*, 70(1), 35-45.
- Anderson, J.A., Colvin, J., & Tobler, N. (2003). Productivity and multi-screen displays. *Technical Report by Lindsay Research & Consulting, Inc.*
- Ankrum, D.R. (1999). Visual ergonomics in the office: Guidelines. *Occupational Health & Safety*, 68(7), 64-74.
- Balakrishnan, R., Fitzmaurice, G., Kurtenbach, G., & Buxton, W. (1999). Digital tape drawing. *Proceedings of the 12th Annual Symposium on User Interface Software and Technology*, 161-169.
- Barfield, W., Lim, R., & Rosenberg, C. (1990). Visual enhancements and geometric field of view as factors in the design of a three-dimensional perspective display. *Proceedings of the 34th Annual Meeting of the Human Factors Society*, 1470-1473.
- Bartram, L., Ware, C., & Calvert, T. (2003). Moticons: Detection, distraction and task. *International Journal of Human-Computer Studies*, 58(5), 515-545.
- Baudisch, P., Good, N., Bellotti, V., & Schraedley, P. (2002). Keeping things in context: A comparative evaluation of focus plus context screens, overviews, and zooming. *Proceedings of the CHI 2002 Conference on Human Factors in Computing Systems*, 259-266.
- Bederson, B.B., Hollan, J.D. (1995). Pad++: A zoomable graphical interface system. *Proceedings of the CHI 1995 Conference on Human Factors in Computing Systems*, 23-24.
- Belton, J. (1992). *Widescreen cinema*. Harvard University Press: Cambridge, MA.
- Bercowicz, D.A., Barnett, G.O., & Chueh, H.C. (1999). eWhiteBoard: A real time clinical scheduler. *American Medical Informatics Association Annual Symposium*.
- Bishop, G., Welch, G. (2000). Working in the Office of "Real Soon Now". *IEEE Computer Graphics and Applications*, 20(4), 76-78.
- Blonder, G. (1996). Graphical passwords. *United States Patent 5559961*.
- Bly, S.A., Rosenberg, J.K. (1986). A comparison of tiled and overlapping windows. *Proceedings of the CHI 1986 Conference on Human Factors in Computing Systems*, 101-106.

- Bobick, A.F., Intille, S.S., Davis, J.W., Baird, F., Pinhanez, C.S., Campbell, L.W., Ivanov, Y.A., Schütte, A., & Wilson, A. (1999). The KidsRoom: A perceptually-based interactive and immersive story environment. *Presence: Teleoperators and Virtual Environments*, 8(4), 367-391.
- Booth, K., Fisher, B., Page, S., Ware, C., & Widen, S. (2000). Wayfinding in a virtual environment. *Graphics Interface 2000*.
- Brignull, H., & Rogers, Y. (2003). Enticing people to interact with large public displays in public spaces. *Proceedings of Interact 2003: IFIP Conference on Human-Computer Interaction*, 17-24.
- Brostoff, S., Sasse, M.A. (2000). Are Passfaces more usable than passwords? A field trial investigation. *Proceedings of HCI on People and Computers XIV*, 405-424.
- Bryson, S., Zeltzer, D., Bolas, M.T., Chapelle, B.d.L., & Bennett, D. (1997). The future of virtual reality: Head-mounted displays versus spatially immersive displays. *Proceedings of SIGGRAPH 1997 International Conference on Computer Graphics and Interactive Techniques*, 485-486.
- Buxton, W. (2001). Less is more (More or less). In Denning, P. (ed.). *The Invisible Future: The seamless integration of technology in everyday life*. McGraw Hill: New York, NY, 145-179.
- Bystrom, K.E., Barfield, W., & Hendrix, C. (1999). A conceptual model of the sense of presence in virtual environments. *Presence: Teleoperators and Virtual Environments*, 8(2), 241-244.
- Cadiz, JJ, Venolia, G.D., Jancke, G., Gupta, A. (2001). *Designing and Deploying an Information Awareness Interface. Proceedings of CSCW 2002*, 314-323.
- Campbell, B.G. (1976). *Humankind emerging*. Little, Brown, & Co: Boston, MA.
- Carpenter, M., Proffitt, D. (2001). Comparing viewer and array mental rotations in different planes. *Memory & Cognition*, 29, 441-448.
- Carrasco, M., Naegele, P. (1995). Visual search as a function of retinal locus of target. *Presented at Annual Meeting of the Psychonomic Society*.
- Chambers, W.W. (1982). Visual technology research simulator: Flight simulation-avionic systems and aero medical aspects. *Royal Aeronautical Society's Flight Simulation Symposium*.
- Chance, S.S., Gaunet, F., Beall, A.C., & Loomis, J.M. (1998). Locomotion mode affects the updating of objects encounter during travel: The contribution of vestibular and proprioceptive inputs to path integration. *Presence: Teleoperators and Virtual Environments*, 7(2), 168-178.
- Chapanis, A., Scarpa, L.C. (1967). Readability of dials at different distances with constant visual angle. *Human Factors*, 9(5), 419-426.

- Churchill, E., Nelson, L., Denoue, L., & Girgensohn, A. (2003). The Plasma Poster network: Posting multimedia content in public places. *Interact 2003: IFIP Conference on Human-Computer Interaction*, 599-606.
- Cook, P., Ellis, C., Graf, M., Rein, G., & Smith, T. (1987). Project Nick: Meetings augmentation and analysis. *ACM Transactions on Information Systems*, 5(2), 132-146.
- Cruz-Neira, C., Sandin, D., & DeFanti, T. (1993). Surround-screen projection-based virtual reality: The design and implementation of the CAVE. *Proceedings of SIGGRAPH 1993 International Conference on Computer Graphics and Interactive Techniques*, 135-142.
- Cutmore, T.R.H., Hine, T.J., Maberly, K.J., Langford, N.M., & Hawgood, G. (2000). Cognitive and gender factors influencing navigation in a virtual environment. *International Journal of Human-Computer Studies*, 53, 223-249.
- Czernuszenko, M., Pape, D., Sandin, D., DeFanti, T., Dawe, G.L., & Brown, M.D. (1997). The ImmersaDesk and Infinity Wall projection-based virtual reality displays. *ACM SIGGRAPH Computer Graphics*, 31(2), 46-49.
- Czerwinski, M., Cutrell, E., & Horvitz, E. (2000). Instant messaging and interruption: influence of task type on performance. *Proceedings of OZCHI 2000*, 356-361.
- Czerwinski, M., Tan, D.S., & Robertson, G.G. (2002). Women take a wider view. *Proceedings of the CHI 2002 Conference on Human Factors in Computing Systems*, 195-202.
- Darken, R., Sibert, J. (1993). A toolset for navigation in virtual environments. *Proceedings of the 6th Annual Symposium on User Interface Software and Technology*, 157-165.
- Darken, R., Sibert, J. (1996). Navigating in large virtual worlds. *The International Journal of Human-Computer Interaction*, 8(1), 49-72.
- Davis, D., Price, W. (1987). *Security for Computer Networks*. Wiley: Chichester, UK.
- de-Brujin, D., de-Mul, S., & van-Oostendorp, H. (1992). The influence of screen size and text layout on the study of text. *Behaviour and Information Technology*, 11(2), 71-78.
- Detenber, B.H., Reeves, B. (1996). A bio-informational theory of emotion: Motion and image size effects on viewers. *Journal of Communication*, 46(3), 66-84.
- Dietz, P., Leigh, D. (2001). DiamondTouch: A multi-user touch technology. *Proceedings of the 14th Annual Symposium on User Interface Software and Technology*, 219-226.
- Dixon, M.W. (1999). Compensation for differences in perceived proportions between large and small displays. Doctoral Dissertation, University of Virginia, Charlottesville, VA.

- Dolezal, H. (1982). *Living in a world transformed: Perceptual and performatory adaptation to visual distortion*. Academic Press: New York, NY.
- Draper, M.H., Viirre, E.S., Furness, T.A., & Gawron, V.J. (2001). Effects of image scale and system time delay on simulator sickness within head-coupled virtual environments. *Human Factors*, 43(1), 129-146.
- Dudfield, H.J., Macklin, C., Fearnley, R., Simpson, A., & Hall, P. (2001). Big is better? Human factors issues of large screen displays with military command teams. *Proceedings of People in Control 2001*, 304-309.
- Duncanson, J.P., Williams, A.D. (1973). *Video conferencing: Reactions of users*. *Human Factors*, 15(5), 471-485.
- Dunn, D., Mikes, N. (2001). *Multiple monitor computing: A significant productivity trend in its very early stage of adoption*, 9X Media Inc.
- Educational Testing Service. (1994). *Practicing to take the GRE test*. New Jersey: ETS.
- Ekstrom, R.B., French, J.W., Harman, H., & Dermen, D. (1976). *Kit of factor-referenced cognitive tests*. Educational Testing Service: Princeton, NJ.
- Elrod, S., Bruce, R., Gold, R., Goldberg, D., Halasz, F., Janssen, W., Lee, D., McCall, K., Pedersen, E., Pier, K., Tang, J., & Welch, B. (1992). Liveboard: A large interactive display supporting group meetings, presentations and remote collaboration. *Proceedings of the CHI 1992 Conference on Human Factors in Computing Systems*, 599-607.
- Epic Games. *Unreal Tournament*. <http://www.unrealtournament.com/>
- European Commission. (2000). *The disappearing computer*. <http://www.cordis.lu/ist/fet/dc.htm>
- Fass, A., Forlizzi, J., & Pausch, R. (2002). MessyDesk and MessyBoard: Two designs inspired by the goal of improving human memory. *Proceedings of Designing Interactive Systems*, 303-311.
- Fish, R.S., Judd, T.H. (1991). A subjective visual quality comparison of NTSC, VHS, and compressed DS1-compatible video. *Proceedings of the Society for Information Display*, 157-164.
- Fitton, D., Cheverst, K. (2003). Experiences managing and maintaining a collection of interactive door displays. *European Symposium on Ambient Intelligence*, 394-409.
- Flach, J. (1990). Control with an eye for perception: Precursors to an active psychophysics. *Ecological Psychology*, 2, 83-111.
- Flesch, R. (1948). A new readability yardstick. *Journal of Applied Psychology*, 32, 221-233.

- Frohlich, D., Thomas, P., Hawley, M., & Hirade, K. (1997). Inaugural issue edition: Future personal computing. *Personal Technologies*, 1, 1-5.
- Fujita, N., Klatzky, R.L., Loomis, J.M., & Golledge, R.G. (1993). The encoding-error model of pathway completion without vision. *Geographical Analysis*, 25, 295-314.
- Furnas, G.W. (1986). Generalized fisheye views. *Proceedings of the CHI 1986 Conference on Human Factors in Computing Systems*, 16-23.
- Gaver, W. (1992). The affordances of media spaces for collaboration. *Proceedings of the 1992 Conference on Computer Supported Cooperative Work*, 17-24.
- Gillie, T. Broadbent, D. (1989). What makes interruptions disruptive? A study of length, similarity, and complexity. *Psychological Research*, 50, 243-50.
- Goldberg, A. (1984). *SMALLTALK-80: The interactive programming environment*. Addison-Wesley Longman Publishing Co., Inc.: Boston, MA.
- Golledge, R.G. (ed.). (1999). *Wayfinding behavior: Cognitive mapping and other spatial processes*. Johns Hopkins University Press: Baltimore, MD.
- Grabe, M.E., Lombard, M., Reich, R.D., Bracken, C.C., & Ditton, T.B. (1999). The role of screen size in viewer experiences of media content. *News Photographer*, 54(4), 4-10.
- Grasso, A., Roulland, F., Snowdon, D., & Muehlenbrock, M. (2002). Designing digital bulletin boards for social networking. *Workshop on Public, Community, and Situated Displays at the 2002 Conference on Computer Supported Cooperative Work*.
- Grayson, D.M. (2000). The role of perceived proximity in video-mediated communication. Doctoral Dissertation, University of Glasgow, Glasgow, Scotland.
- Greenberg, S., Rounding, M. (2001). The notification collage: posting information to public and personal displays. *Proceedings of the CHI 2001 Conference on Human Factors in Computing Systems*, 514-521.
- Grudin, J. (2001). Partitioning digital worlds: Focal and peripheral awareness in multiple monitor use. *Proceedings of the CHI 2001 Conference on Human Factors in Computing Systems*, 458-465.
- Guilford, J.P. (1972). Thurstone's primary mental abilities and structure-of-intellect abilities. *Psychological Bulletin*, 77(2), 129-143.
- Guilford, J.P., Zimmerman, W.S. (1948). The Guilford-Zimmerman aptitude survey. *Journal of Applied Psychology*, 32, 24-34.
- Guimbretière, F. (2002). Fluid interaction for high resolution wall-size displays. Doctoral Dissertation, Stanford University, Stanford, CA.

- Gutekunst, T., Bauer, D., Caronni, G., Hasan, & Plattner, B. (1995). A distributed and policy-free general purpose shared window system. *IEEE/ACM Transactions on Networking*.
- Hagen, M.A., Jones, R.K., & Reed, E.S. (1978). On a neglected variable in theories of pictorial perception: Truncation of the visual field. *Perception & Psychophysics*, 23(4), 326-330.
- Halevi, S., Krawczyk, H. (1999). Public-key cryptography and password protocols. *ACM Transactions on Information and System Security*, 2(3), 230-168.
- Hall, E. (1966). *The hidden dimension*. Doubleday: Garden City, NY.
- Haller, N., Metz, C., Nesser, P., Straw M. (1998). A one-time password system. *Network Working Group Request for Comments* 2289.
- Hatada, T., Sakata, H., & Kusaka, H. (1980). Psychophysical analysis of the perception of reality induced by a widefield display. *SMPTE Journal*, 89, 560-569.
- Henderson, A., Card, S.K. (1986). Rooms: The use of multiple virtual workspaces to reduce space contention in a window-based graphical user interface. *ACM Transactions on Graphics*, 5(3), 211-243.
- Hess, S.M., Detweiler, M.C., & Ellis, R.D. (1999). The utility of display space in keeping track of rapidly changing information. *Human Factors*, 41, 257-281.
- Hiniker, P.J. (1998). The common operational picture as evolving schema for command centers as complex adaptive systems. *Proceedings of the 1998 Symposium on Command and Control Research*.
- Hiniker, P.J., Entin, E.E. (1992). Examining cognitive processes in command crisis: New HEAT experiments on shared battle graphics and time tagging. *Proceedings of the 1992 Symposium on Command and Control Research*.
- Hitchings, J. (1995). Deficiencies of the traditional approach to information security and the requirements for a new methodology. *Computers and Security*, 14, 377-383.
- Hollan, J., Hutchins, E., & Kirsh, D. (2000). Distributed cognition: Toward a new foundation for human-computer interaction research. *ACM Transactions on Computer-Human Interaction*, 7(2), 174-196.
- Hollan, J., Stornetta, S. (1992). Beyond being there. *Proceedings of the CHI 1992 Conference on Human Factors in Computing Systems*, 119-125.
- Hosman, R., van der Haart, J. (1981). Effects of vestibular and visual motion perception on task performance. *Acta Psychologica*, 48, 271-281.
- Humphreys, G., Hanrahan, P. (1999). A distributed graphics system for large tiled displays. *Proceedings of IEEE Visualization 1999*, 215-224.

- Husemann, D. (ed.). (2000). Pervasive computing (special issue). *Computer Networks*, 35(4).
- Hutchings, D.G., Stasko, J. (2002). New operations for display space management and window management. *Georgia Institute of Technology Technical Report GIT-GVU-02-18*.
- ID Software. *Doom, Quake*. <http://www.idsoftware.com/>
- Infield, S.E. (1991). An investigation into the relationship between navigation skill and spatial abilities. Doctoral Dissertation, University of Washington, Seattle, WA. Dissertation Abstracts International, 52(5-B), 2800.
- Intel Image Processing Library. For more info, see:
<http://developer.intel.com/software/products/perflib/ipl/>
- Ishii, H., & Ullmer, B. (1997). Tangible bits: Towards seamless interfaces between people, bits and atoms. *Proceedings of the CHI 1997 Conference on Human Factors in Computing Systems*, 234-241.
- Izadi, S., Brignull, H., Rodden, T., Rogers, Y., & Underwood, M. (2003). Dynamo: A public interactive surface supporting the cooperative sharing and exchange of media. *Proceedings of the 16th Annual Symposium on User Interface Software and Technology*, 159-168.
- Jackson, L.A. (1992) *Physical appearance and gender*. Albany, NY: State University of New York Press.
- Jain, A., Hong, L., Pankanti, S. (2000). Biometric identification. *Communications of the ACM*, 43(2), 90-98.
- Jaschinski-Kruza, W. (1990). On the preferred viewing distances to screen and document at VDU workplaces. *Ergonomics*, 33(8), 1055-1063.
- Jermyn, I., Mayer, A., Monroe, F., Reiter, M.K., Rubin, A. (1999). The design and analysis of graphical passwords. *8th USENIX Security Symposium*.
- Jesty, L.C. (1958). The relation between picture size, viewing distance, and picture quality. *Proceedings of the IEE*, 105B, 425-439.
- Johanson, B., Fox, A., & Winograd, T. (2002). The interactive workspaces project: experiences with ubiquitous computing rooms. *IEEE Pervasive Computing Magazine*, 1(2), 67-74.
- Josephs, R.A, Geisler, R.B., & Silvera, D.H. (1994). Judgment by quantity. *Journal of Experimental Psychology*, 123, 21-32.
- Just, M. A., Carpenter, P. A. (1985). Cognitive coordinate systems: Accounts of mental rotation and individual differences in spatial ability. *Psychological Review*, 92(2), 137-172.

- Kandogan, E., Schneiderman, B. (1997). Elastic windows: Evaluation of multi-window operations. *Proceedings of the CHI 1997 Conference on Human Factors in Computing Systems*, 250-257.
- Kaufman, J.E., Christensen, J.F. (1987). *IES Lighting Handbook: 1987 Application Volume*. Illuminating Engineering Society of North America: New York, NY.
- Kearns, M.J., Warren, W.H., Duchon, A.P., & Tarr, M.J. (2002). Path integration from optic flow and body senses in a homing task. *Perception*, 31, 349-374.
- Kenyon, R.V., Kneller, E.W. (1993). The effects of field of view size on the control of roll motion. *IEEE Transactions on Systems, Man, and Cybernetics*, 23(1), 183-193.
- Kim, T. (1996). The memroy and persuasion effects of presence in television advertising processes. Doctoral Dissertation, University of North Carolina, Chapel Hill, NC.
- Kirsh, D. (1993). The intelligent use of space. *Artificial Intelligence*, 73, 31-68.
- Klatzky, R.L., Loomis, J.M., Beall, A.C., Chance, S.S., & Golledge, R.G. (1998). Spatial updating of self-position and orientation during real, imagined, and virtual locomotion. *Psychological Science*, 9(4), 293-298.
- Kreifeldt, J.G. & McCarthy, M.E. (1981). Interruption as a test of the user-computer interface. *Proceedings of the 17th Annual Conference on Manual Control*, 655-667.
- Kucera, H., Francis, W.N. (1967). *Computational analysis of present-day American English*. Brown University Press: Providence, RI.
- Lai, J., Levas, A., Chou, P., Pinhanetz, C., & Viveros, M. (2002). BlueSpace: Personalizing workspace through awareness and adaptability. *International Journal of Human-Computer Studies*, 57(5), 415-428.
- Lang, P.J. (1980). Behavioral treatmeent and bio-behavioral treatment: Computer applications. In Sidowski, J.B., Johnson, J.H., & Williams, T.A. (eds.), *Technology in Mental Health Care Delivery Systems*, 119-137.
- Lang, P.J. (1995). The network model of emotion: Motivational connections. In Wyer, R.S., & Srull, T.K. (eds.), *Advances in Social Cognition*.
- Leibe, B., Starner, T., Ribarsky, W., Wartell, Z., Krum, D., Weeks, J., Singletary, B., & Hodges, L. (2000). Towards spontaneous interaction with the perceptive workbench, a semi-immersive virtual environment. *IEEE Computer Graphics and Applications*, 20(6), 54-65.
- Levenshtein, V. (1966). Binary codes capable of correcting deletions, insertions and reversals. *Soviet Physics Daklady*, 10, 707-710.
- Li, D., Li, R. (2002). Transparent sharing and interoperation of heterogeneous single-user applications. *Proceedings of the 2002 Conference on Computer Supported Cooperative Work*, 246-255.

- Li, K., Chen, H., Chen, Y., Clark, D., Cook, P., Damianakis, S., Essl, G., Finkelstein, A., Funkhouser, T., Housel, T., Klein, A., Liu, Z., Praun, E., Singh, J., Shedd, B., Pal, J., Tzanetakis, G., & Zheng, J. (2000). Building and using a scalable display wall system. *IEEE Computer Graphics and Applications*, 20(4), 29-37.
- Lim, R.W., Wogalter, M.S. (2000). The position of static and on-off banners in WWW displays on subsequent recognition. *Proceedings of the IEA 2000/HFES 2000 Congress*, 420-423.
- Lohman, D. (1979). Spatial ability: A review and reanalysis of the correlational literature, *Technical Report, N.8. Stanford University, Aptitude Research Project, School of Education*.
- Lombard, M. (1995). Direct responses to people on the screen: Television and personal space. *Communication Research*, 22(3), 288-324.
- Lombard, M., Ditton, T.B., Grabe, M.E., & Reich, R.D. (1997). The role of screen size in viewer responses to television fare. *Communication Reports*, 10(1), 95-106.
- Lombard, M., Reich, R.D., Grabe, M.E., Bracken, C.C., & Ditton, T.B. (2000). Presence and television: The role of screen size. *Human Communication Research*, 26(1), 75-98.
- Loomis, J.M., Klatzky, R.L., Golledge, R.G., Cicinelli, J.G., Pellegrino, J.W., & Fry, P.A. (1993). Nonvisual navigation by blind and sighted: Assessment of path integration ability. *Journal of Experimental Psychology: General*, 122(1), 73-91.
- Lund, A.M. (1993). The influence of video image size and resolution on viewing-distance preferences. *SMPTE Journal*, 102, 406-415.
- MacIntyre, B., Cowan, W.B. (1992). A practical approach to calculating luminance contrast on a CRT. *ACM Transactions on Graphics*, 11(4), 336-347.
- MacIntyre, B., Mynatt, E.D., Volda, S., Hansen, K.M., Tullio, J., & Corso, G.M. (2001). Support for multitasking and background awareness using interactive peripheral displays. *Proceedings of the 14th Annual Symposium on User Interface Software and Technology*, 41-50.
- Maglio, P.P., Campbell, C.S. (2000). Tradeoffs in displaying peripheral information. *Proceedings of the CHI 2000 Conference on Human Factors in Computing Systems*, 241-248.
- Mandryk, R., Scott, S., & Inkpen, K. (2002). Display factors influencing co-located collaboration. *Proceedings of the 2002 Conference on Computer Supported Cooperative Work*.
- Mantei, M. (1988). Capturing the capture lab concepts: A case study in the design of computer supported meeting environments. *Proceedings of the 1988 Conference on Computer Supported Cooperative Work*, 257-270.
- Mark, G. (2002). Extreme collaboration. *Communications of the ACM*, 45(6), 89-93.

- Matsushita, N., Rekimoto, J. (1997). HoloWall: Designing a finger, hand, body, and object sensitive wall. *Proceedings of the 10th Annual Symposium on User Interface Software and Technology*, 209-210.
- Maxivista. <http://www.maxivista.com/>
- McCarthy, J.F., Costa, T.J., & Liongosari, E.S. (2001). UniCast, OutCast & GroupCast: An exploration of new interaction paradigms for ubiquitous, peripheral displays. *Proceedings of 3rd International Conference on Ubiquitous Computing*, 332-345.
- McCrickard, D., Catrambone, R., Stasko, J. (2001) Evaluating animation in the periphery as a mechanism for maintaining awareness. *Proceedings of Interact 2001: IFIP Conference on Human-Computer Interaction*, 148-152.
- McFarlane, D.C. & Latorella, K.A. (2002). The scope and importance of human interruption in HCI design. *Human-Computer Interaction*, 17, 1-62.
- Microsoft Corporation. *Remote Desktop*.
<http://www.microsoft.com/windowsxp/downloads/tools/rdclientdl.msp>
- Milner, A.D., Goodale, M.A. (1996). *The visual brain in action*. Oxford, UK: Oxford University Press.
- Moran, T.P., Saund, E., van Melle, W., Gujar, A., Fishkin, K.P., & Harrison, B.L. (1999). Design and technology for Collaborage: Collaborative collages of information on physical walls. *Proceedings of the 12th Annual Symposium on User Interface Software and Technology*, 197 - 206.
- MSN Instant Messenger. <http://messenger.msn.com>
- Myers, B.A. (1988). A taxonomy of window manager user interfaces. *IEEE Computer Graphics and Applications*, 8(5), 65-84.
- Mynatt, E.D., Igarashi, T., Edwards, W.K., & LaMarca, A. (1999). Flatland: New dimensions in office whiteboards. *Proceedings of the CHI 1999 Conference on Human Factors in Computing Systems*, 346-353.
- Nakagawa, M., Oguni, T., Yoshino, T., Horiba, K., & Sawada, S. (1996). Interactive dynamic whiteboard for educational applications. *International Conference on Virtual Systems and Multimedia 1996*, 479-484.
- Nathan, J.G., Anderson, D.R., Field, D.E., & Collins, P. (1985). Television viewing at home: Distances and visual angles of children and adults. *Human Factors*, 27(4), 467-476.
- National Computation Science Alliance. *Display Wall-In-A-Box*.
<http://www.ncsa.uiuc.edu/Projects/AllProjects/Projects83.html>.
- National Institute of Standards and Technology. (1995). Password usage. *NIST FIPS PUB 112*.

- Neuman, W.R. (1990) Beyond HDTV: Exploring subjective responses to very high definition television. *Research Report for GTE Labs and the TVOT Consortium*. MIT, Cambridge, MA.
- Neumann, P.G. (1994). Risks of passwords. *Communications of the ACM*, 37(4), 126.
- New York Times on the Web: <http://www.nytimes.com>
- Newsweek (1992). 9 March, 8.
- Norman, N. (1998). *The invisible computer*. MIT Press: Cambridge, MA.
- O'Hara, K., Perry, M., Churchill, E., & Russell, D. (2003). *Public and situated displays: Social and interactional aspects of shared display technologies*. Kluwer Academic Publishers: Dordrecht.
- Ohtani, T., Mitshuhashi, T. (1973). Picture quality of a high definition television system. *TEBS*, 3-4.
- Patrick, E., Cosgrove, D., Slavkovic, A., Rode, J.A., Verratti, T., & Chiselko, G. (2000). Using a large projection screen as an alternative to head-mounted displays for virtual environments. *Proceedings of the CHI 2000 Conference on Human Factors in Computing Systems*, 478-485.
- Pedersen, E., McCall, K., Moran, T.P., & Halasz, F. (1993). Tivoli: An electronic whiteboard for informal workgroup meetings. *Proceedings of the CHI 1993 Conference on Human Factors in Computing Systems*, 391-398.
- Perlin, K., Fox, D. (1993). Pad: An alternative approach to the computer interface. *Computer graphics proceedings, Annual Conference Series*, 57-64.
- Péruch, P., May, M., & Wartenberg, F. (1997). Homing in virtual environments: Effects of field of view and path layout. *Perception*, 26, 301-311.
- Philbeck, J.W., Klatzky, R.L., Behrmann, M., Loomis, J.M., & Goodridge, J. (2001). Active control of locomotion facilitates nonvisual navigation. *Journal of Experimental Psychology: Human Perception and Performance*, 27, 141-153.
- Piantanida, T.P., Boman, D., Larimer, J., Gille, J., & Reed, C. (1992). Studies of the field-of-view/resolution tradeoff in virtual reality. *Proceedings of the SPIE: Human Vision, Visual Processing and Digital Display III*, 448-456.
- PowerWall. University of Minnesota.
<http://www.lcse.umn.edu/research/powerwall/powerwall.html>.
- Rantanen, E.M., Goldberg, J.H. (1999). The effect of mental workload on the visual field size and shape. *Ergonomics*, 42(6), 816-834.
- Raskar, R., van Baar, J., Willwacher, T., & Rao, S. (2004). Quadric transfer for immersive curved screen displays. *Proceedings of Eurographics 2004*, 23(3).

- Raskar, R., Welch, G., Cutts, M., Lake, A., Stesin, L., & Fuchs, H. (1998). The Office of the Future: A unified approach to image-based modeling and spatially immersive displays. *Proceedings of SIGGRAPH 1998 International Conference on Computer Graphics and Interactive Techniques*, 179-188.
- RealVNC. *Virtual Network Computing*. <http://www.realvnc.com/>
- Reeves, B., Detenber, B., & Steuer, K. (1993). New televisions: The effects of big pictures and big sound on viewer responses to the screen. *Annual Conference of the International Communication Association*.
- Reeves, B., Lombard, M., & Melwani, G. (1992). Faces on the screen: Pictures or natural experience? *Annual Conference of the International Communication Association*.
- Reeves, B., Nass, C. (1996). *The Media Equation: How people treat computers, television, and new media like real people and places*, 193-293.
- Reeves, B., Thorson, E. (1986). Watching television: Experiments on the viewing process. *Communication Research*, 13(3), 343-361.
- Rekimoto, J. (1998). A multiple device approach for supporting whiteboard-based interactions. *Proceedings of SIGGRAPH 1998 International Conference on Computer Graphics and Interactive Techniques*, 344-351.
- Rekimoto, J., Saitoh, M. (1999). Augmented Surfaces: A Spatially Continuous Work Space for Hybrid Computing Environments. *Proceedings of SIGGRAPH 1999 International Conference on Computer Graphics and Interactive Techniques*, 378-385.
- Ribchester, E. (1958). Discussion before the radio and telecommunication section. *Proceedings of the IEE*, 105B, 437.
- Richardson, T., Quentin, S., Wood, K.R., & Hopper, A. (1998). Virtual Network Computing. *IEEE Internet Computing*, 2(1), 33-38.
- Riecke, B.E., van Veen, H.A.H.C., & Bülthoff, H.H. (2000). Visual homing is possible without landmarks: A path integration study in virtual reality. *Max-Planck-Institut für biologische Kybernetik, Germany*.
- Rieser, J. J. (1989). Access to knowledge of spatial structure at novel points of observation. *Journal of Experimental Psychology: Learning, Memory and Cognition*, 15, 1157-1165.
- Ruddle, R., Payne, S., & Jones, D. (1999). The effects of maps on navigation and search strategies in very-large-scale virtual environments. *Journal of Experimental Psychology: Applied*, 5, 54-75.
- Russell, D.M., Trimble, J., & Wales, R. (2002). Two paths from the same place: Task driven and human-centered evolution of a group interaction surface. *Make IT Easy Conference*.

- Ryer, A. (1997). *Light measurement handbook*. <http://www.intl-light.com/handbook/>.
- Sadashige, K., Ando, S. (1984). Recent development in large screen video display equipment technology. In *Television Image Quality*. SMPTE: Scarsdale, NY.
- Sandor, P.B., Leger, A. (1991). Tracking with a restricted field of view. *Aviation, Space, and Environmental Medicine*, 62(11), 1026-1031.
- Schacter, D.L. (1987). Implicit memory: History and current status. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 13(3), 501-518.
- Schell, J., Shochet, J. (2001). Designing interactive theme park rides. *IEEE Computer Graphics and Applications*, 21(4), 11-13.
- Schikore, D.R., Fischer, R.A., Frank, R., Gaunt, R., Hobson, J., & Whitlock, B. (2000). High-resolution multi-projector display walls and applications. *IEEE Computer Graphics and Applications*, 20(4), 38-44.
- Schmidt, A. (2000). Implicit human computer interaction through context. *Personal Technologies*, 4, 191-199.
- ShareKMC. <http://www.whatisnew.com/article966.html>
- Shepard, R.N., Metzler, J. (1971). Mental rotation of three-dimensional objects. *Science*, 171(3972), 701-703.
- Shoemaker, G.B.D. (2001). Single Display Groupware research in the year 2000. *Technical Report TR 2001-1*. Simon Fraser University.
- Siegel, A., White, S. (1975). The development of spatial representations of large-scale environments. In H. Reese (ed.), *Advances in Child Development and Behavior*, 10, New York: Academic Press, 10-55.
- Silvera, D.H., Josephs, R.A., & Giesler, R.B. (2002). Bigger is better: the influence of physical size on aesthetic preference judgements. *Journal of Behavioral Decision Making*, 15, 189-202.
- Simmons, T. (2001). What's the optimum computer display size? *Ergonomics in Design*, 9(4), 19-24.
- Smart Technologies, Inc. *SMART Board*. <http://www.smarttech.com/>.
- Smith, D.C., Irby, C., Kimball, R., Verplank, B., & Harslem, E. (1982). Designing the Star user interface. *Byte*, 242-282.
- Smith, G., Baudisch, P., Robertson, G.G., Czerwinski, M., Meyers, B., Robbins, D., Horvitz, E., & Andrews, D. (2003). Groupbar: The taskbar evolved. *Proceedings of OZCHI 2003: New Directions in Interaction*, 34-43.
- Spence, R., Apperly, M. (1982). Data base navigation: an office environment for the professional. *Behaviour and Information Technology*, 1(1), 43-54.

- Stefik, M., Foster, G., Bobrow, D., Kahn, K., Lanning, S., & Suchman, L. (1987). Beyond the chalkboard: Computer support for collaboration and problem solving in meetings. *Communications of the ACM*, 30(1), 32-47.
- Stewart, J., Bederson, B.B., & Druin, A. (1999). Single Display Groupware: A model for co-present collaboration. *Proceedings of the CHI 1999 Conference on Human Factors in Computing Systems*, 286-293.
- Stille, S., Minocha, S., Ernst, R. (1997). An adaptive window management system. *Proceedings of Interact 1997: IFIP Conference on Human-Computer Interaction*, 67-68.
- Streitz, N.A., Geißler, J., Holmer, T., Konomi, S., Müller-Tomfelde, C., Reisch, W., Rexroth, P., Seitz, P., & Steinmetz, R. (1999). i-LAND: An interactive landscape for creativity and innovation. *Proceedings of the CHI 1999 Conference on Human Factors in Computing Systems*, 120-127.
- Streitz, N.A., Röcker, C., Prante, T., Stenzel, R., & van Alphen, D (2003). Situated interaction with ambient information: facilitating awareness and communication in ubiquitous work environments. In Harris, D., Duffy, V., Smith, M., & Stephanidis, C. (eds.), *Human-Centered Computing: Cognitive, Social, and Ergonomic Aspects*. Lawrence Erlbaum Publishers: New Jersey, 133-137.
- Sukthankar, R., Cham, T.J., Sukthankar, G. (2001). Dynamic shadow elimination for multi-projector displays. *Proceedings of IEEE Conference on Computer Vision and Pattern Recognition 2001*.
- Summet, J., Abowd, G.D., Corso, G.M., & Rehg, J.M. (2003). Virtual rear projection: A comparison study of projection technologies for large interactive displays. *Technical Report GIT-GVU-03-36*, Georgia Institute of Technology, Atlanta, GA.
- Sutherland, I. (1963). Sketchpad. Doctoral Dissertation, Massachusetts Institute of Technology, Boston, MA.
- Swaminathan, K., Sato, S. (1997). Interaction design for large displays. *Interactions*, 4(1), 15-24.
- Swanson, T.L., Couvillion, W. (2001). Display systems for VR-based training: When is seeing believing? *Interservice/Industry Training Systems and Education Conference*.
- Synergy. <http://sourceforge.net/projects/synergy2/>.
- Tan, D.S., Czerwinski, M. (2003a). Effects of visual separation and physical discontinuities when distributing information across multiple displays. *Proceedings of OZCHI 2003*.
- Tan, D.S., Czerwinski, M. (2003b). Information voyeurism: Social impact of physically large displays on information privacy. *Extended Abstracts of the CHI 2003 Conference on Human Factors in Computing Systems*, 748-749.

- Tan, D.S., Czerwinski, M., Robertson, G.G. (2003). Women Go With the (Optical) Flow. *Proceedings of the CHI 2003 Conference on Human Factors in Computing Systems*, 209-215.
- Tan, D.S., Gergle, D., Scupelli, P., & Pausch, R. (2003). With similar visual angles, larger displays improve performance on spatial tasks. *Proceedings of the CHI 2003 Conference on Human Factors in Computing Systems*, 217-224.
- Tan, D.S., Gergle, D., Scupelli, P., & Pausch, R. (2004). Physically large displays improve path integration in 3D virtual navigation tasks. *Proceedings of the CHI 2004 Conference on Human Factors in Computing Systems*, 439-446.
- Tan, D.S., Meyers, B., & Czerwinski, M. (2004). WinCuts: Manipulating arbitrary window regions for more effective use of screen space. *Extended Abstracts of the CHI 2004 Conference on Human Factors in Computing Systems*, 1525-1528.
- Tan, D.S., Pausch, R. (2002). Pre-emptive Shadows: Eliminating the blinding light from projectors. *Extended Abstracts of the CHI 2002 Conference on Human Factors in Computing Systems*, 682-683.
- Tan, D.S., Stefanucci, J.K., Proffitt, D.R., & Pausch, R. (2001). The Infocockpit: Providing location and place to aid human memory. *Proceedings of the 2001 Workshop on Perceptive User Interfaces*, 1-4.
- Tani, M., Horita, M., Yamaashi, K., & Tanikoshi, K. (1994). Courtyard: Integrating shared overview on a large screen and per-user detail on individual screens. *Proceedings of the CHI 1994 Conference on Human Factors in Computing Systems*, 44-50.
- Thomas, P., Gellersen, H. (2001). Ensemble computing. *International Journal of Human-Computer Interaction*, 13(2), 113-121.
- Thorndyke, P., Hayes-Roth, B. (1982). Differences in spatial knowledge acquired from maps and navigation. *Cognitive Psychology*, 14, 560-589.
- Thurstone, L.L., Thurstone, T.G. (1941). Factorial studies of intelligence. *Psychometric Monograph*, 2.
- Tjan, B. (1996). Color spaces for human observer. *Technical Note, Minnesota Laboratory for Low-Vision Research*, University of Minnesota.
- Touchdown Entertainment. *Jupiter Game Development Platform*.
<http://www.touchdownentertainment.com/>
- Trimble, J., Wales, R., & Gossweiler, R. (2002). NASA position paper for the cscw 2002 workshop on public, community and situated displays: MERBoard. *Workshop on Public, Community, and Situated Displays at the 2002 Conference on Computer Supported Cooperative Work*.
- Verdac, A.N. (1968). *Stage to screen*. Benjamin Blom: New York, NY.

- Waller, D., Hunt, E., & Knapp, D. (1998). The transfer of spatial knowledge in virtual environment training. *Presence: Teleoperators and Virtual Environments*, 7(2), 129-143.
- Weiser, M. (1991) The computer for the 21st century. *Scientific American*, 265(3), 66-75.
- Weiser, M., Brown, J.S. (1996). *The coming age of calm technology*. Palo Alto, CA: Xerox PARC.
- Welch, G., Fuchs, H., Raskar, R., Towles, H., & Brown, M.S. (2000). Projected imagery in your "Office of the Future". *IEEE Computer Graphics and Applications*, 20(4), 62-67.
- Wells, M.J., Venturino, M. (1990). Performance and head movements using a helmet-mounted display with different sized fields-of-view. *Optical Engineering*, 29(8), 870-877.
- Westerink, J., Roufs, J. (1989). Subjective image quality as a function of viewing distance, resolution, and picture size. *SMPTE Journal*, 98, 113-119.
- Wickens, C.D., Carswell, C.M. (1995). The proximity compatibility principle: Its psychological foundation and relevance to display design. *Human Factors*, 37, 473-494.
- Wickens, C.D., Hollands, J. (2000). *Engineering psychology and human performance*, 3rd edition. Prentice Hall: New Jersey, 69-118.
- Wisneski, G., Ishii, H., Dahley, A., Gorbett, M., Brave, S., Ullmer, B., & Yarin, P. (1998). Ambient display: Turning architectural space into an interface between people and digital information. In N.A. Streitz, S. Konomi, & H.J. Curkhardt (eds.), *Cooperative buildings*. Springer-Verlag: Heidelberg, Germany.
- Witmer, B., Bailey, J., Knerr, B., & Parsons, K. (1996). Virtual spaces and real world places: Transfer of route knowledge. *International Journal of Human-Computer Studies*, 45, 413-428.
- Wolfe, J.M., O'Neill, P., & Bennett, S.C. (1998). Why are there eccentricity effects in visual search? Visual and attentional hypotheses. *Perception and Psychophysics*, 60(1), 140-156.
- X Windows. <http://www.xwin.net/>
- Zviran, M., Haga, W.J. (1990). Cognitive passwords: The key to easy access control. *Computers and Security*, 9(8), 723-736.